

# Geophysical Research Letters®



## RESEARCH LETTER

10.1029/2025GL120472

### Key Points:

- We propose a probabilistic air-sea turbulent momentum and heat flux algorithm based on neural networks trained on in situ observations
- Our algorithm quantifies the uncertainty (variability) around the mean, with a mean prediction similar to existing bulk algorithms
- Marked seasonal sensitivities are found in deterministic and stochastic tests performed on forced single-column upper-ocean model GOTM

### Supporting Information:

Supporting Information may be found in the online version of this article.

### Correspondence to:

J. Wu,  
[jiarong.wu@nyu.edu](mailto:jiarong.wu@nyu.edu)

### Citation:

Wu, J., Perezhagin, P., Gagne, D. J., Reichl, B. G., Subramanian, A. C., Thompson, E. J., & Zanna, L. (2026). Data-driven probabilistic air-sea flux parameterization. *Geophysical Research Letters*, 53, e2025GL120472. <https://doi.org/10.1029/2025GL120472>

Received 8 NOV 2025

Accepted 4 MAR 2026

### Author Contributions:

**Conceptualization:** Jiarong Wu, David John Gagne, Aneesh C. Subramanian, Laure Zanna

**Data curation:** Aneesh C. Subramanian, Elizabeth J. Thompson

**Formal analysis:** Jiarong Wu

**Funding acquisition:** David John Gagne, Brandon G. Reichl, Laure Zanna

**Investigation:** Jiarong Wu

**Methodology:** Jiarong Wu, Pavel Perezhagin, Brandon G. Reichl, Laure Zanna

**Resources:** Laure Zanna

**Software:** Jiarong Wu, Pavel Perezhagin

**Supervision:** David John Gagne, Laure Zanna

© 2026. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

## Data-Driven Probabilistic Air-Sea Flux Parameterization

Jiarong Wu<sup>1</sup> , Pavel Perezhagin<sup>1</sup> , David John Gagne<sup>2</sup> , Brandon G. Reichl<sup>3</sup> , Aneesh C. Subramanian<sup>4</sup> , Elizabeth J. Thompson<sup>5</sup> , and Laure Zanna<sup>1</sup> 

<sup>1</sup>Courant Institute of Mathematical Sciences, New York University, New York, NY, USA, <sup>2</sup>National Center for Atmospheric Research, Boulder, CO, USA, <sup>3</sup>NOAA—Geophysical Fluids Dynamics Laboratory, Princeton, NJ, USA, <sup>4</sup>Department of Atmospheric and Oceanic Sciences, University of Colorado Boulder, Boulder, CO, USA, <sup>5</sup>NOAA Physical Sciences Lab, Boulder, CO, USA

**Abstract** Accurately quantifying air-sea fluxes is important for understanding air-sea interactions and improving coupled weather and climate models. This study introduces a probabilistic framework to represent the highly variable nature of air-sea fluxes, which is missing in deterministic bulk algorithms. Assuming Gaussian distributions conditioned on the input variables, we use artificial neural networks and eddy-covariance measurement data to estimate the mean and variance by minimizing negative log-likelihood loss. The trained neural networks provide alternative mean flux estimates to existing bulk algorithms, and quantify the uncertainty around the mean estimates. A stochastic parameterization of air-sea turbulent fluxes can be constructed by sampling from the predicted distributions. Tests in a single-column forced upper-ocean model suggest that changes in flux algorithms influence sea surface temperature and mixed layer depth seasonally. The ensemble spread in stochastic runs is most pronounced during spring restratification.

**Plain Language Summary** Understanding the exchange of heat and momentum between the ocean and the atmosphere is key to improving weather and climate predictions. However, quantifying these exchanges (fluxes) is difficult, and models rely on simplified equations to compute the flux values based on state variables (wind speed, air temperature, sea surface temperature, etc.). This study introduces a new statistical method based on machine learning (artificial neural networks) trained using direct field observations of fluxes to better estimate these ocean-atmosphere fluxes. The model provides not only a better estimate of the average flux, but also an understanding of its uncertainty. Preliminary tests show that this new flux estimation method has a considerable effect on the simulated state of the upper ocean, especially during certain seasonal shifts. This approach helps improve the accuracy of air-sea flux estimates and can eventually lead to better coupled weather and climate models.

## 1. Introduction

The atmosphere and ocean exchange mass, momentum, and heat across the air-sea interface. These fluxes influence ocean and atmosphere processes across a vast range of scales. However, quantifying air-sea fluxes is challenging, and significant uncertainties remain (Cronin et al., 2019).

In this work, we focus on momentum flux (denoted as  $\tau_x$  and  $\tau_y$ ), and on turbulent heat flux (THF), which is a non-radiative heat flux consisting of sensible heat flux ( $Q_S$ ) due to air-sea temperature difference and latent heat flux ( $Q_L$ ) due to air-sea humidity difference. Direct in situ observations of these turbulent fluxes from atmosphere to ocean rely on measuring the covariances of turbulent fluctuations:

$$\tau_x = -\rho_a \overline{u'w'}, \tau_y = -\rho_a \overline{v'w'}, Q_S = -\rho_a c_p \overline{w'T'}, Q_L = -\rho_a L_e \overline{w'q'}. \quad (1)$$

Here  $u'$ ,  $v'$ ,  $w'$ ,  $T'$ , and  $q'$  are fluctuations of three velocity components, potential temperature, and specific humidity;  $\rho_a$  is the air density;  $c_p$  is the specific heat capacity at constant pressure, and  $L_e$  is the latent heat of evaporation. The measurements are taken in the atmospheric surface layer, where these fluxes are assumed to be constant (Fairall et al., 1996).

The challenges of quantifying these fluxes lie in both observational and modeling aspects. Due to the requirements of sophisticated instruments and careful quality control, such direct measurements are usually carried out on designated research cruises (Bradley & Fairall, 2006). Less equipped measuring platforms and remote sensing measure atmospheric and oceanic surface variables (wind speed, temperature, humidity, etc.) and rely on

**Validation:** Jiarong Wu, Brandon G. Reichl  
**Visualization:** Jiarong Wu  
**Writing – original draft:** Jiarong Wu, Pavel Perezhogin, David John Gagne, Brandon G. Reichl, Aneesh C. Subramanian, Elizabeth J. Thompson, Laure Zanna

a bulk flux algorithm to compute the fluxes from the mean observed variables. The same algorithm is used in coupled weather and climate models to compute fluxes based on prognostic state variables of the oceanic and atmospheric surfaces.

The widely used flux algorithms are termed bulk algorithms because they use bulk quantities from the surface layer to model surface fluxes. The momentum and turbulent heat fluxes are formulated as proportional to the magnitude of wind speed  $|U_a|$  and air-sea difference in velocity, temperature, and humidity

$$\tau_x = \rho_a C_D S (U_a - U_o), \quad Q_S = \rho_a c_p C_H S (T_a - T_o), \quad Q_L = \rho_a L_e C_E S (q_a - q_s), \quad (2)$$


where  $C_D$ ,  $C_H$ , and  $C_E$  are the transfer coefficients for momentum, sensible, and latent heat, respectively.  $S$  is the scalar wind speed relative to the ocean surface (subject to gustiness correction). Here,  $\tau_x$  is aligned with the surface wind and the cross-wind component  $\tau_y$  is assumed to be zero. The transfer coefficients are calculated based on the Monin-Obukhov similarity theory, where certain parameters (stability function and roughness length) are empirically determined from observations; see a complete description in, for example, Fairall et al. (2003) for COARE algorithm.

Currently, there are different bulk algorithms, fitted to different sets of observations (Biri et al., 2023; Brunke et al., 2003). Varying levels of simplifications and empirical corrections (e.g., cool-skin warm-layer or gustiness corrections) also exist. This source of uncertainty affects flux estimation in general circulation models (GCMs) and flux products. Sensitivity studies have shown that changes in bulk algorithms can considerably affect atmospheric dynamics (Harrop et al., 2018; Hsu et al., 2022; Polichtchouk & Shepherd, 2016) and oceanic state (Bonino et al., 2022). Flux products also suffer from uncertainty in bulk algorithms, in addition to uncertainty in bulk inputs (Yu, 2019). Due to the prohibitive computational cost of high-fidelity numerical simulations, eddy-covariance (EC) measurements remain our best “ground-truth” for calibrating bulk algorithms, despite their sparsity and intrinsic measurement uncertainty (Gleckler & Weare, 1997). In this study, we use a quality-controlled research vessel cruise data set provided by NOAA Physical Sciences Lab (PSL) to develop an alternative data-driven flux algorithm.

Another potential drawback of deterministic bulk algorithms is the lack of variability. Bulk algorithms are simplified representations of the underlying dynamical processes and, at best, represent a statistically averaged value of fluxes given the observables. There is a significant spread of EC flux data around the prediction of bulk algorithms. Such deviations from the mean flux values may be crucial for modeling processes on small and fast scales (Nuijens et al., 2024); they can also have a rectifying effect for large-scale dynamics in nonlinear GCMs. More complex formulas such as sea-state-dependent parameterization (Bouin et al., 2024; Edson et al., 2013) have been proposed that incorporate additional sources of variability, but have not yet been fully validated and widely adopted.

We aim to quantify the uncertainty (variability) around deterministic flux algorithms in EC data while taking an agnostic perspective on its sources. This is achieved by applying a framework based on conditional parametric probability distributions (Barnes et al., 2021; Guillaumin & Zanna, 2021; Nix & Weigend, 1994; Schreck et al., 2024). In particular, we use artificial neural networks (ANNs) to approximate the distribution parameters. Recent advances in machine learning have led to its applications in surface flux parameterization (Bourras et al., 2007; Cummins et al., 2024; Leufen & Schädler, 2019; McCandless et al., 2022; Zhou et al., 2024), although mainly for point estimates. In most of the paper, we use the terms “uncertainty” and “variability” interchangeably to refer to the spread in observed data, but its attribution will be discussed in the end.

An estimate of uncertainty around deterministic bulk algorithms will benefit uncertainty quantification in downstream applications such as flux products and stochastic air-sea flux parameterization. Stochastic parameterizations have been shown to improve the representation of low-frequency variability and to reduce bias in the mean state (Andrejczuk et al., 2016; Berner et al., 2017; Juricke et al., 2017; Williams, 2012). In particular, air-sea fluxes are expected to play a significant role in the upper ocean. However, the sensitivity of the upper ocean states to uncertainty in the flux algorithm is not yet well quantified. We implement our probabilistic flux algorithm in the single-column General Ocean Turbulence Model (GOTM, Umlauf and Burchard (2005)) to study the effects on the upper ocean states.



**Figure 1.** (a) Ship trajectories of the various cruises in the NOAA PSL data set. (b) An illustration of the ANN-based conditional Gaussian probabilistic model. Note that we are visualizing only two of the input space dimensions, for the purpose of showing the concept of a conditional Gaussian distribution. (c) Distributions of input variables for different subsets of data.

## 2. Data and Model

### 2.1. Direct In Situ Covariance Measurements

The air-sea flux observation data set we use is collected by decades of research cruises conducted by NOAA PSL. There are about 10,000 samples after quality control, and all are ship-borne, hourly-averaged covariance measurements according to Equation 1. The trajectories of the various cruises are shown in Figure 1a. Among the labeled cruises, Metz is a compilation of multiple earlier cruises in the 1990s; WHOTS, EPIC, DYNAMO are conducted in the tropics and Stratus in the subtropics; Calwater, NEAQS, HIWInGS in mid/high latitudes of the northern hemisphere, and Capricorn, GasEX in the southern ocean. Overall, this data set covers various geographical locations, although a disproportional amount of data is collected in the tropics (more than 50%). We note that analyzing the measurements taken along transects as time series will likely reveal additional embedded information, but in this study we treat them as independent samples of the underlying conditional distribution.

### 2.2. Probabilistic Model

We consider observed bulk variables  $\mathbf{X}$  (such as wind speed, humidity, etc.), unobserved variables  $\mathbf{Z}$  (e.g., sea state, vertical wind profile, etc.), and targeted flux outputs  $\mathbf{Y} = (\tau_x, \tau_y, Q_s, Q_L)$ . Due to the unobserved variables  $\mathbf{Z}$ , the relation  $\mathbf{Y} = \mathbf{F}_{model}(\mathbf{X})$  is not deterministic even when the physical laws governing the fluxes  $\mathbf{Y} = \mathbf{F}_{true}(\mathbf{X}, \mathbf{Z})$  are deterministic. Using the proposed probabilistic model, we account for all uncertainties in the data, and further attributions are discussed in Section 5.

We assume that each flux component  $y$  of  $\mathbf{Y}$  follows a uni-variate conditional Gaussian distribution with mean  $\mu(\mathbf{X})$  and standard deviation (std)  $\sigma(\mathbf{X})$ :

$$y \sim \mathcal{N}(\mu(\mathbf{X}), \sigma^2(\mathbf{X})). \quad (3)$$

In other words,  $\mu(\mathbf{X})$  is our best unbiased estimation of the flux component  $y$  given  $\mathbf{X}$ , and the errors are distributed according to  $\mathcal{N}(0, \sigma^2(\mathbf{X}))$  when conditioned on  $\mathbf{X}$ . This idea of conditional Gaussian distribution is demonstrated in Figure 1b with the latent heat flux data.

Having chosen the parametric distribution, we aim to learn the parameters  $\mu_\theta(\mathbf{X})$  and  $\sigma_\phi^2(\mathbf{X})$  with the given data. Here  $\theta$  and  $\phi$  denote the learnable parameters in the data-driven models for  $\mu$  and  $\sigma^2$ . The optimization procedure

is based on computing the probability as a function of parameters (i.e., likelihood) of observing a sample value  $y$  given the sample input features  $\mathbf{x}$ :

$$p(y|\mathbf{x}, \theta, \phi) = \frac{1}{\sqrt{2\pi\sigma_\phi^2(\mathbf{x})}} \exp \left[ -\frac{(y - \mu_\theta(\mathbf{x}))^2}{2\sigma_\phi^2(\mathbf{x})} \right]. \quad (4)$$

The optimal model parameters  $\theta$  and  $\phi$  minimize this negative log-likelihood loss, computed on the training data set summed over a total of  $N_{\text{sample}}$  (Nix & Weigend, 1994):

$$L_{\text{nl}}(\theta, \phi) = \sum_{m=1}^{N_{\text{sample}}} \frac{1}{2} \left[ \log(\sigma_\phi^2(\mathbf{x}_m)) + \frac{(y_m - \mu_\theta(\mathbf{x}_m))^2}{\sigma_\phi^2(\mathbf{x}_m)} \right] + \text{const.} \quad (5)$$

### 2.3. ANN Architecture and Training

The mathematical framework described in Section 2.2 is general and applies to any data-driven model that approximates parametric distributions. In particular, as shown in Figure 1b, we use two ANNs to represent the mean  $\mu_\theta$  and the variance  $\sigma_\phi^2$  for each flux component. In this case,  $\theta$  and  $\phi$  are the weights and biases of the ANNs. The nonlinear activation function of the hidden layers is the Sigmoid function. The positivity of the predicted variance  $\sigma_\phi^2$  is ensured with the exponential activation function in the last ANN layer.

The input variables are selected based on a balance between expressibility and the risk of overfitting, as well as practical considerations for the typically available variables in GCMs. After testing different numbers of inputs and their combinations, we have chosen five inputs

$$\mathbf{X} = (U_a, T_a, T_o, RH, p_a) \quad (6)$$

which are wind speed (measured between 16 and 21 m), atmospheric surface temperature (measured between 12 and 19.5 m), sea surface temperature (measured by sea snake at 0.05 m depth), relative humidity (measured between 12 and 19.5 m), and atmospheric pressure at sea surface. Using relative humidity instead of specific humidity gives better behaviors during training, although the two can be converted given atmospheric temperature and pressure. We do not include the heights at which the meteorological variables are measured, although they are used as inputs to existing bulk algorithms. This is mainly because the ship-borne sensors are placed at a height where the vertical gradient of the atmospheric surface variables is vanishing and the height dependence cannot be explicitly learned. This is not a problem for climate-scale GCMs, which have a coarse vertical resolution, but it is a potential caveat when applied to low-lying platforms, such as buoys.


The training of ANNs is described in detail in Supporting Information S1. Briefly summarized here, we first train the networks on mean-square error loss and then on negative log-likelihood loss (Equation 5). This two-step procedure allows the mean ANN to capture more variability in the data. The small data limit and the distribution shift shown in Figure 1c are challenging for data-driven models, which we overcome through input feature selection, model design, training strategies (such as early stopping), and cross-validation.

Having learned the parametric distribution of the fluxes for the given inputs, predictions can be made in two ways: using the mean  $\mu_\theta$  for deterministic predictions or sampling from the distribution  $\mathcal{N}(\mu_\theta, \sigma_\phi^2)$  for stochastic predictions. In Section 3.1, we evaluate the statistical scores of the deterministic predictions. In Section 3.2, we analyze the predicted dependence of  $\mu_\theta$  and  $\sigma_\phi^2$  on inputs. In Section 4, numerical experiments using both deterministic and stochastic heat fluxes are discussed. For simplicity, we denote the mean predictions as  $(\mu_{\tau_x}, \mu_{\tau_y}, \mu_{Q_s}, \mu_{Q_L})$  and the variance predictions as  $(\sigma_{\tau_x}^2, \sigma_{\tau_y}^2, \sigma_{Q_s}^2, \sigma_{Q_L}^2)$ , omitting subscript  $\theta$  and  $\phi$ .

## 3. Evaluation of the ANN-Based Probabilistic Air-Sea Flux Model

### 3.1. Statistical Scores of the Deterministic Predictions

We evaluate the statistical scores of the flux predictions in terms of the root-mean-square-error RMSE, the coefficient of determination  $R^2$ , and bias of the estimate (compared to observations):



**Figure 2.** (a) ANN-based deterministic predictions (orange dots) and bulk algorithm predictions (black dots) plotted against measured fluxes. We are only visualizing 10% of total samples. (b)  $R^2$  and bias of ANN (orange) and bulk algorithm (black), evaluated on different geographical subsets. (c) ANN prediction of the mean of  $Q_s$  on a uniform input grid. Among the five inputs, we fix sea surface temperature  $T_o = 10^\circ\text{C}$ , relative humidity  $RH = 80\%$ , and sea level pressure  $p_a = 1,010$  hPa. The dashed gray lines mark zero temperature difference  $T_a = T_o$ . (d) ANN prediction of the std of  $Q_s$  for the same grid. (e) Scatter plot of all data points and their marginal distribution. Gaussian kernel density estimation provides the gray contour lines that indicate the available samples per unit  $\Delta U$  and  $\Delta(T_a - T_o)$ . The same contour lines are overlaid on (c, d).

$$\text{RMSE}(\hat{y}, y) = (\mathbb{E}[(\hat{y} - y)^2])^{1/2}, \quad (7)$$

$$R^2(\hat{y}, y) = 1 - \mathbb{E}[(\hat{y} - y)^2] / \text{Var}[y], \quad (8)$$

$$\text{Bias}(\hat{y}, y) = \mathbb{E}[\hat{y} - y] \quad (9)$$

Here  $\hat{y}$  is the algorithm prediction and  $y$  is the truth, in our case the EC measurement ( $\tau_{x,c}$ ,  $\tau_{y,c}$ ,  $Q_{s,c}$ ,  $Q_{L,c}$ ). We evaluate these metrics for the ANN-based deterministic prediction ( $\mu_{\tau_x}$ ,  $\mu_{\tau_y}$ ,  $\mu_{Q_s}$ ,  $\mu_{Q_L}$ ) compared to a baseline bulk algorithm ( $\tau_{x,b}$ , 0,  $Q_{s,b}$ ,  $Q_{L,b}$ ). In particular, we choose COARE 3.6 (Fairall et al., 2003), which is the most up-to-date version of COARE. It is worth mentioning that COARE was fitted to the same shipborne EC data we used, with additional cruises and buoy measurements.

Figure 2a shows the flux predictions plotted against the EC measurements. ANN-based deterministic predictions have similar but slightly higher  $R^2$  than COARE for all three fluxes, see legends in Figure 2a, with only a subset of the input variables. For the cross-wind momentum flux  $\tau_y$  (not shown here), ANN-based deterministic prediction has little predictive skills, as there is no input variable to indicate the sign of cross-wind stress. Without additional



input features (such as surface wave information), it is reasonable to accept zero prediction as is in bulk algorithms. However, we can quantify the variability around the zero prediction with the current framework.

The statistical scores of both ANN and COARE vary considerably between different turbulent fluxes. The sensible heat flux  $Q_S$  is the hardest to predict for both based on its lowest  $R^2$  values, while the latent heat flux  $Q_L$  gives the largest magnitude of error in terms of RMSE. Another observation is that scores of both ANN and COARE also vary considerably between geographical locations. Figure 2b shows  $R^2$  and bias of three fluxes evaluated on different subsets of cruises grouped roughly by geographical locations (same as in Figure 1). This is because each geographical subset has substantially different distributions of the input variable, as shown in Figure 1c. In particular,  $R^2$  for sensible heat flux  $Q_S$  is very low in the Tropics, only around 0.15. This is partly because the variance of the sensible heat flux is small there. Overall, the fluctuation of predictive scores over different regions is shared between ANN and COARE.

The most salient improvement of ANN compared to baseline bulk algorithm is for latent heat flux, which is generally the least constrained of all turbulent fluxes. In certain regions, ANN sees up to around 0.2 increase in  $R^2$ . Overall, we notice that the estimate of latent heat given by bulk algorithm seems biased (toward predicting more evaporation from the ocean), and ANN fixes the bias to some extent, as shown in Figure 2b. This bias correction in latent heat flux was also found in a recent paper using data-driven method (Zhou et al., 2024) and may have important implications for global simulations. We comment that a latent heat flux model accounting for sea spray generation may offer more flexibility and help mitigate this bias (Andreas et al., 2015; Barr et al., 2023). Table S1 in Supporting Information S1 lists RMSE and  $R^2$  evaluated on the full data set and on different geographical subsets. The reason underlying the statistical improvement is likely that the flexible structure of ANN provides more expressibility, and therefore fits the data better where certain assumptions of traditional bulk algorithms break down, for example, the logarithmic profile for humidity and the expression for scalar roughness lengths. This idea is further discussed in Supporting Information S1 and ANN predictions in terms of transfer coefficients  $C_D$ ,  $C_H$ , and  $C_E$  are also shown (Figure S1 in Supporting Information S1).

### 3.2. Structure of the Mean and Variance Predicted by ANNs

In addition to the statistical scores, we further evaluate the ANN predictions on a reduced-dimensional uniform grid to probe and interpret the ANN flux model. This is especially important for understanding how the predicted std depends on input variables. Figures 2c and 2d illustrate this using sensible heat flux as an example.

Figure 2c shows the prediction of  $Q_S$  mean by ANN. Unlike bulk algorithms that assume  $Q_{S,b} \propto S(T_a - T_o)$ , the structure of ANN is more flexible, allowing non-zero flux values even when  $T_a = T_o$ . A similar figure for COARE is shown in Figure S2 in Supporting Information S1. We did not impose any constraints on the sign of  $Q_S$ , contrary to Zhou et al. (2024), which used a penalty in the loss function to ensure matching signs of  $T_a - T_o$  and  $Q_S$ . In our experience, such a penalty is not necessary for a good statistical fit. Furthermore, there is insufficient evidence for a strictly down-gradient assumption, since locally counter-gradient fluxes are possible (Blay-Carreras et al., 2014; Deardorff, 1972) and systematic bias in measurements may exist.

Figure 2d shows the prediction of std of  $Q_S$  for the same grid. Uncertainty increases with wind speed, as expected, since residuals should scale with flux magnitude to some extent. A small but finite uncertainty persists at very low wind speeds, likely due to measurement challenges associated with small fluxes. For a given wind speed, the uncertainty is consistently higher for  $T_a < T_o$ , corresponding to unstable boundary layer conditions. As shown in Figure 2e, data are sparse for wind speeds above 15 m/s and large temperature differences. While the ANN seems to extrapolate reasonably well, it is worth exploring methods to further constrain it where training samples are not available. We note that the same issue applies to existing bulk algorithms as well.

## 4. Idealized Sensitivity Tests of Probabilistic Air-Sea Flux Parameterization

To assess the impact of the proposed stochastic air-sea flux parameterization, we implement it in the single column model GOTM (Umlauf & Burchard, 2005), which provides a controlled environment without the complex interactions with other nonlinear processes present in GCMs and significantly reduces the computational cost of running ensembles. The governing equations are a set of 1D diffusion-type equations

$$\partial_t u = -\partial_z \overline{w' u'} + f v, \quad \partial_t v = -\partial_z \overline{w' v'} - f u, \quad \partial_t T = -\partial_z \overline{w' T'}, \quad \partial_t S = -\partial_z \overline{w' S'}, \quad (10)$$

where the prognostic variables are the horizontal velocity components  $u$  and  $v$  (subject to Coriolis force), temperature  $T$ , and salinity  $S$ . The turbulent fluxes  $\overline{w' u'}$ ,  $\overline{w' v'}$ ,  $\overline{w' T'}$ ,  $\overline{w' S'}$  are not resolved but are instead parameterized by vertical mixing schemes. For robustness, we test our air-sea flux algorithms in combination with two ocean surface boundary layer (OSBL) mixing parameterizations: the K-profile-parameterization (KPP) scheme (Large et al., 1994) and a more sophisticated second-order closure  $k - \epsilon$  scheme (Umlauf & Burchard, 2005), both used in GCMs but typically under different resolutions. Air-sea momentum and heat fluxes serve as boundary conditions to Equation 10. Additionally, surface fluxes can impact the mixing parameterization.

#### 4.1. THF Time Series at OWS Papa

We use Ocean Weather Station (OWS) Papa (50°N, 145°W) as our test site. This long-term North Pacific station provides hourly meteorological and oceanic measurements, including temperature and salinity profiles. Its weak horizontal advection makes it a common benchmark for testing vertical mixing parameterization in single-column models (Li et al., 2021; Umlauf & Burchard, 2005). Here, our goal is simply to use this location to demonstrate the use of stochastic flux parameterization and to assess the potential impact on upper ocean state, with given vertical mixing parameterizations. We discuss the flux predictions given the observed input variables at OWS Papa before showing results from the forced experiments.

Figure 3a shows an example of ANN-predicted (deterministic) THF time series  $Q = \mu_{Q_s} + \mu_{Q_L}$ , compared to bulk algorithms, zoomed in to July and October 2015. We plot two bulk algorithms—COARE 3.6 and NCAR (Large & Yeager, 2009)—to show that the ANN deviations exceed those between different bulk algorithms. As with the THF itself, the difference between the ANN and bulk predictions is seasonal, since the input distributions vary throughout the year. Figure 3c shows the monthly averaged flux discrepancy between ANN and COARE 3.6. The months of July and October see the most negative (more heat flux out of the ocean) and positive differences. Notice that within a period of high THF around October 10, ANN predicts the least extreme negative values while NCAR predicts the most extreme negative value, reflecting all three algorithms' divergence at high wind speed, where few measurements exist to constrain them.


Figure 3a also shows the ANN-predicted  $\pm 1 \sigma_Q$  for THF. We assume no covariance between the sensible and latent heat flux residuals, so the uncertainty of their sum simply follows  $\sigma_Q^2 = \sigma_{Q_s}^2 + \sigma_{Q_L}^2$ . Importantly, the predicted  $\sigma_Q$  is state dependent, so the uncertainty envelope changes over time. Given the ANN-predicted mean and std, we can create an ensemble of stochastically perturbed fluxes, and two examples are shown in Figure 3a. The stochastic perturbation is expected to stem from certain unobserved processes with characteristic correlation time scales. Therefore, instead of white noise, we use temporally correlated noise  $\epsilon$  generated through a discrete autoregressive model of order one—AR(1) (Giglio et al., 2017). The correlation time of 60 hr is estimated from the OWS Papa data. Note that the bulk algorithm predicted THF is generally within  $\pm \sigma_Q$ , and thus can be interpreted as one possible ensemble member.

Despite relatively weak horizontal advection, there is still a net heat flux imbalance at OWS Papa (Large et al., 1994) if horizontal advection is ignored. We mitigate this by restarting the simulation with observed temperature and salinity profiles every month and by emphasizing inter-comparison between simulations rather than direct validation against observations. Details of the numerical experiments are provided in Supporting Information S1.

#### 4.2. Sensitivity of the Upper Ocean State to THF Uncertainty

With the monthly restart and stochastically perturbed THF ensemble, our tests can be understood as experiments that quantify the short-term uncertainty in the upper ocean state due to uncertainty in the flux algorithm. In the following discussion, we focus on two characteristic quantities: sea surface temperature (SST) and mixed layer depth (MLD).

At OWS Papa, the mixed layer is deep in the winter because strong winds and surface cooling induce convective mixing, and shallow in the summer because weaker winds and more radiative heating lead to stable stratification. In spring, as the radiative heating gains and the THF decreases in magnitude, there is a fast restratification process,



**Figure 3.** (a) Example THF time series for July and October 2015, computed using observed input variables at OWS Papa. Black solid line: COARE 3.6; black dashed: NCAR; orange line: ANN. Orange shade shows one standard deviation for THF predicted by ANN, which is used to generate noise-perturbed heat fluxes (examples shown by thin orange lines). (b) Difference of temperature profiles in the deterministic single-column run  $\Delta T(z, t) = T_{ANN}(z, t) - T_{bulk}(z, t)$ . Black and orange dashed lines show the diagnosed MLD for bulk and ANN, respectively. (c) Monthly-averaged THF discrepancies between ANN and COARE. (d) SST and (e) MLD responses in the deterministic runs between (ANN minus COARE). Trends with inconsistent signs over the years are masked with hatching. (f) Monthly-averaged SST and (g) monthly-averaged MLD spread in 20 ensemble runs. The spread is monthly averaged while the instantaneous value can be larger (see Figure S5 in Supporting Information S1).

whereas in fall, the erosion of the stratification is more gradual. This annual cycle is relatively well captured given the monthly restart, as shown in Figure S3 in Supporting Information S1.

Related to this annual cycle, our main finding is a significant seasonality in the sensitivity to surface heat flux. Figure 3b shows the difference in temperature profiles between ANN and COARE flux forced runs, as well as the MLD, for 2015. There is a perceivable difference in temperature (fairly uniform across the mixed layer) as a result of a change in surface heat flux. ANN generally predicts more surface heat loss, which deepens the mixed layer and cools the sea surface. When the mixed layer is the shallowest in the summer months, there is the largest change in SST, consistent with a heat budget idea. The seasonality is confirmed by the composite statistics for 4 years of monthly-averaged change in SST and MLD are summarized in Figures 3d and 3e. We note that these effects are consistent across  $k - \epsilon$  and KPP (see Figure S6 in Supporting Information S1), suggesting that the effects of air-sea fluxes is somewhat insensitive to the OSBL parameterization used.

Perhaps more important than comparison between any two specific deterministic runs, the ensemble of stochastically perturbed simulations provides an estimate of the variability in SST and MLD induced by the variability of THF. From the 20 ensemble member runs, we do not observe systematic drift induced by the stochastic



forcing; the ensemble mean of the stochastic runs closely tracks the deterministic run (see Figure S5 in Supporting Information S1). This is expected since we are performing forced experiments and there is no feedback of state variables (such as SST) to surface forcing. However, we observe considerable spread between ensemble members, which is summarized in Figures 3f and 3g, again with clear seasonal patterns. The largest spread in MLD occurs during the months of April and May, since stochasticity can significantly affect the onset of restratification. This may have important implications for biogeochemical processes, suggesting that comparisons between modeled and observed upper ocean profiles, especially during restratification, may be sensitive to surface THF uncertainty. The largest spread in SST is seen around August, again because the mixed layer is shallow and SST is sensitive to surface heat fluxes. As in deterministic runs, the effects seen in stochastic runs are not particularly sensitive to different OSBL parameterizations.

## 5. Conclusion and Discussion

We propose an ANN-based probabilistic model for turbulent air-sea fluxes. State-of-the-art bulk algorithms represent the statistical mean of eddy-covariance flux observations for a given set of state variables. Our work reinforces this idea by providing an alternative mean estimate using a purely data-driven approach. The ANN-predicted mean is similar to existing bulk algorithms, with marginally higher statistical correlation to data and reduced bias for latent heat flux. Additionally, the ANN-predicted standard deviation around the mean offers a measure of uncertainty (variability). Tests using OWS Papa data and a single-column model show that the difference in THF between the different algorithms' estimates exhibits signs of seasonality, as does the sensitivity of upper ocean states, despite relatively small magnitudes. Stochastic ensemble runs indicate that the spread of MLD is largest during spring restratification, while the spread of SST is largest during summer when the mixed layer is shallow.

We comment that the current framework can also be used to estimate the standard deviation around a given mean estimation, for example, any existing bulk algorithm, by considering the residual between the observation and the mean estimation. It is important to note that the uncertainty we estimate does not distinguish between measurement uncertainty (both instrument and sampling, with sampling uncertainty being the dominant factor, see e.g., Rannik et al., 2016) and missing physics (the inability to account for all predictive input variables). Such a partition is difficult and likely depends on the averaging window of the EC flux, which will be examined in future studies. Nonetheless, the ANN-estimated uncertainty highlights regions in the input space where the data is more "scattered." This motivates field campaigns for better sampling of these regimes and examination of measurement uncertainty, as well as the use of high-fidelity numerical simulation to constrain the flux estimate (Clayson et al., 2023).

There are several possible extensions of this work in the future. The ANN model uses a limited set of input variables, which is why the improvement over the existing algorithm may be considered marginal. Further improvement is possible, as several additional variables are measured and available in the data set (e.g., turbulent kinetic energy), but many of them are not yet prognostic variables in GCMs. The height dependence is difficult to learn from the current data set alone, but incorporating data from other measurement platforms (e.g., buoy) might help. We also need more high-quality vertical profile measurements of surface wind, temperature, and humidity. The probabilistic model presented assumes a univariate conditional Gaussian distribution for each flux component, which may not fully capture the underlying complexity of unobserved physical processes. The model can be extended to relax the Gaussian assumption and include the covariance between flux components. It would also be valuable to explore potential connections between the data-driven formulation and ad hoc modifications commonly used in flux algorithms, such as the gustiness factor. Finally, the analysis at OWS Papa provides an initial estimate of the sensitivity magnitude but is unlikely to be representative of the global ocean with varying meteo-ocean conditions and coupling dynamics across regions. We plan to extend our analysis to global simulations to assess the broader impact of the proposed air-sea flux algorithm.

In summary, the probabilistic framework proposed in this paper provides a formal way to represent uncertainty in air-sea flux estimates, which is particularly important given the extreme scarcity of direct observations, despite their continued role as the primary means for constraining air-sea flux parameterizations. While existing bulk algorithms represent the mean values of turbulent fluxes given limited input variables, our approach takes a step toward examining the variability around the mean values. Stochastic air-sea flux parameterization offers a promising alternative to the deterministic approach. It is crucial to prescribe the appropriate magnitude and

correlation scale of noise with varying model resolutions, which needs to be better understood in future studies. In the single-column test, the spread induced by stochastic residuals exceeds that of different deterministic flux algorithms. Therefore, extending the testing to global GCMs could potentially affect the variability and address long-standing biases through interaction with horizontal transport and other nonlinear dynamics.

## Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

## Availability Statement

The NOAA PSL data are documented and available here <https://downloads.psl.noaa.gov/psd3/cruises/>. The code for ANN training and evaluation, as well as the model weights are openly available (Wu, 2025). The GOTM code is available here <https://gotm.net/portfolio/>. In particular, this work used an implementation of GOTM in the PDE solver Basilisk (Popinet, 2020) [http://basilisk.fr/src/test/ows\\_papa.c](http://basilisk.fr/src/test/ows_papa.c). For comparison to bulk algorithms, we used the aerobulk-python package <https://github.com/xgcm/aerobulk-python>.

## Acknowledgments

We acknowledge funding from NSF through the Learning the Earth with Artificial Intelligence and Physics (LEAP) Science and Technology Center (STC) (Award #2019625). We thank Julia Simpson, Pierre Gentine, and Bia Villas Bôas for helpful discussions and the anonymous reviewers for their useful feedback. Aakash Sane, Qing Li, and Stéphane Popinet provided valuable advice regarding the GOTM implementation. We acknowledge the continued efforts from NOAA PSL lab (including Ludovic Bariteau, Chris Fairall, and many others) in obtaining and publishing the EC data sets. The statements, findings, conclusions, and recommendations are those of the author(s) and do not necessarily reflect the views of the National Oceanic and Atmospheric Administration, or the U.S. Department of Commerce. This research was also supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

## References

- Andreas, E. L., Mahrt, L., & Vickers, D. (2015). An improved bulk air–sea surface flux algorithm, including spray-mediated transfer. *Quarterly Journal of the Royal Meteorological Society*, 141(687), 642–654. <https://doi.org/10.1002/qj.2424>
- Andrejczuk, M., Cooper, F. C., Juricke, S., Palmer, T. N., Weisheimer, A., & Zanna, L. (2016). Oceanic stochastic parameterizations in a seasonal forecast system. *Monthly Weather Review*, 144(5), 1867–1875. <https://doi.org/10.1175/MWR-D-15-0245.1>
- Barnes, E. A., Barnes, R. J., & Gordillo, N. (2021). Adding uncertainty to neural network regression tasks in the geosciences. *arXiv*. <https://doi.org/10.48550/arXiv.2109.07250>
- Barr, B. W., Chen, S. S., & Fairall, C. W. (2023). Sea-state-dependent sea spray and air–sea heat fluxes in tropical cyclones: A new parameterization for fully coupled atmosphere–wave–ocean models. <https://doi.org/10.1175/JAS-D-22-0126.1>
- Berner, J., Achatz, U., Batté, L., Bengtsson, L., Cámara, A. D. L., Christensen, H. M., et al. (2017). Stochastic parameterization: Toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3), 565–588. <https://doi.org/10.1175/BAMS-D-15-00268.1>
- Biri, S., Cornes, R. C., Berry, D. I., Kent, E. C., & Yelland, M. J. (2023). AirSeaFluxCode: Open-source software for calculating turbulent air-sea fluxes from meteorological parameters. *Frontiers in Marine Science*, 9, 1049168. <https://doi.org/10.3389/fmars.2022.1049168>
- Blay-Carreras, E., Pardyjak, E. R., Pino, D., Alexander, D. C., Lohou, F., & Lathon, M. (2014). Countergradient heat flux observations during the evening transition period. *Atmospheric Chemistry and Physics*, 14(17), 9077–9085. <https://doi.org/10.5194/acp-14-9077-2014>
- Bonino, G., Iovino, D., Brodeau, L., & Masina, S. (2022). The bulk parameterizations of turbulent air–sea fluxes in NEMO4: The origin of sea surface temperature differences in a global model study. *Geoscientific Model Development*, 15(17), 6873–6889. <https://doi.org/10.5194/gmd-15-6873-2022>
- Bouin, M.-N., Lebeaupin Brossier, C., Malardel, S., Voldoire, A., & Sauvage, C. (2024). The wave-age-dependent stress parameterisation (WASP) for momentum and heat turbulent fluxes at sea in SURFEX v8.1. *Geoscientific Model Development*, 17(1), 117–141. <https://doi.org/10.5194/gmd-17-117-2024>
- Bourras, D., Reverdin, G., Caniaux, G., & Belamari, S. (2007). A nonlinear statistical model of turbulent air–sea fluxes. *Monthly Weather Review*, 135(3), 1077–1089. <https://doi.org/10.1175/MWR3335.1>
- Bradley, F., & Fairall, C. (2006). *A guide to making climate quality meteorological and flux measurements at sea* (Tech. Rep.). NOAA Technical Memorandum OAR PSD.
- Brunke, M. A., Fairall, C. W., Zeng, X., Eymard, L., & Curry, J. A. (2003). Which bulk aerodynamic algorithms are least problematic in computing ocean surface turbulent fluxes? *Journal of Climate*.
- Clayson, C. A., DeMott, C., De Szoeke, S., Chang, P., Foltz, G., Krishnamurthy, R., et al. (2023). *A new paradigm for observing and modeling of air-sea interactions to advance Earth system prediction* (Tech. Rep.). U.S. CLIVAR Project Office. <https://doi.org/10.5065/24j7-w583>
- Cronin, M. F., Gentemann, C. L., Edson, J., Ueki, I., Bourassa, M., Brown, S., et al. (2019). Air-sea fluxes with a focus on heat and momentum. *Frontiers in Marine Science*, 6, 430. <https://doi.org/10.3389/fmars.2019.00430>
- Cummins, D. P., Guemas, V., Blein, S., Brooks, I. M., Renfrew, I. A., Elvidge, A. D., & Prytherch, J. (2024). Reducing parametrization errors for polar surface turbulent fluxes using machine learning. *Boundary-Layer Meteorology*, 190(3), 13. <https://doi.org/10.1007/s10546-023-00852-8>
- Deardorff, J. W. (1972). Numerical investigation of neutral and unstable planetary boundary layers. *Journal of the Atmospheric Sciences*.
- Edson, J. B., Jampana, V., Weller, R. A., Bigorre, S. P., Plueddemann, A. J., Fairall, C. W., et al. (2013). On the exchange of momentum over the open ocean. *Journal of Physical Oceanography*, 43(8), 1589–1610. <https://doi.org/10.1175/JPO-D-12-0173.1>
- Fairall, C. W., Bradley, E. F., Hare, J. E., Grachev, A. A., & Edson, J. B. (2003). Bulk parameterization of air–sea fluxes: Updates and verification for the COARE algorithm. *Journal of Climate*.
- Fairall, C. W., Bradley, E. F., Rogers, D. P., Edson, J. B., & Young, G. S. (1996). Bulk parameterization of air-sea fluxes for tropical ocean-global atmosphere coupled-ocean atmosphere response experiment. *Journal of Geophysical Research*, 101(C2), 3747–3764. <https://doi.org/10.1029/95JC03205>
- Giglio, D., Gille, S. T., Subramanian, A. C., & Nguyen, S. (2017). The role of wind gusts in upper ocean diurnal variability. *Journal of Geophysical Research: Oceans*, 122(9), 7751–7764. <https://doi.org/10.1002/2017JC012794>
- Gleckler, P. J., & Weare, B. C. (1997). Uncertainties in global ocean surface heat flux climatologies derived from ship observations. *Journal of Climate*, 10(11), 2764–2781. [https://doi.org/10.1175/1520-0442\(1997\)010<2764:UIGOSH>2.0.CO;2](https://doi.org/10.1175/1520-0442(1997)010<2764:UIGOSH>2.0.CO;2)
- Guillaumin, A. P., & Zanna, L. (2021). Stochastic-deep learning parameterization of ocean momentum forcing. *Journal of Advances in Modeling Earth Systems*, 13(9), e2021MS002534. <https://doi.org/10.1029/2021MS002534>

- Harrop, B. E., Ma, P.-L., Rasch, P. J., Neale, R. B., & Hannay, C. (2018). The role of convective gustiness in reducing seasonal precipitation biases in the tropical West Pacific. *Journal of Advances in Modeling Earth Systems*, 10(4), 961–970. <https://doi.org/10.1002/2017MS001157>
- Hsu, C.-W., DeMott, C. A., Branson, M. D., Reeves Eyre, J., & Zeng, X. (2022). Ocean surface flux algorithm effects on tropical Indo-Pacific intraseasonal precipitation. *Geophysical Research Letters*, 49(7), e2021GL096968. <https://doi.org/10.1029/2021GL096968>
- Juricke, S., Palmer, T. N., & Zanna, L. (2017). Stochastic subgrid-scale ocean mixing: Impacts on low-frequency variability. *Journal of Climate*, 30(13), 4997–5019. <https://doi.org/10.1175/JCLI-D-16-0539.1>
- Large, W. G., McWilliams, J. C., & Doney, S. C. (1994). Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameterization. *Reviews of Geophysics*, 32(4), 363–403. <https://doi.org/10.1029/94RG01872>
- Large, W. G., & Yeager, S. G. (2009). The global climatology of an interannually varying air–sea flux data set. *Climate Dynamics*, 33(2), 341–364. <https://doi.org/10.1007/s00382-008-0441-3>
- Leufen, L. H., & Schädler, G. (2019). Calculating the turbulent fluxes in the atmospheric surface layer with neural networks. *Geoscientific Model Development*, 12(5), 2033–2047. <https://doi.org/10.5194/gmd-12-2033-2019>
- Li, Q., Bruggeman, J., Burchard, H., Klingbeil, K., Umlauf, L., & Bolding, K. (2021). Integrating CVMix into GOTM (v6.0): A consistent framework for testing, comparing, and applying ocean mixing schemes. *Geoscientific Model Development*, 14(7), 4261–4282. <https://doi.org/10.5194/gmd-14-4261-2021>
- McCandless, T., Gagne, D. J., Kosović, B., Haupt, S. E., Yang, B., Becker, C., & Schreck, J. (2022). Machine learning for improving surface-layer-flux estimates. *Boundary-Layer Meteorology*, 185(2), 199–228. <https://doi.org/10.1007/s10546-022-00727-4>
- Nix, D., & Weigend, A. (1994). Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)* (Vol. 1, pp. 55–60). <https://doi.org/10.1109/ICNN.1994.374138>
- Nuijens, L., Wenegrat, J., Lopez Dekker, P., Pasquero, C., O'Neill, L. W., Arduin, F., et al. (2024). The air–sea interaction (ASI) submesoscale: Physics and impact. In *Lorentz Workshop*.
- Polichtchouk, I., & Shepherd, T. G. (2016). Zonal-mean circulation response to reduced air–sea momentum roughness. *Quarterly Journal of the Royal Meteorological Society*, 142(700), 2611–2622. <https://doi.org/10.1002/qj.2850>
- Popinet, S. (2020). A vertically-Lagrangian, non-hydrostatic, multilayer model for multiscale free-surface flows. *Journal of Computational Physics*, 418, 109609. <https://doi.org/10.1016/j.jcp.2020.109609>
- Rannik, I., Peltola, O., & Mammarella, I. (2016). Random uncertainties of flux measurements by the eddy covariance technique. *Atmospheric Measurement Techniques*, 9(10), 5163–5181. <https://doi.org/10.5194/amt-9-5163-2016>
- Schreck, J. S., Gagne, D. J., Becker, C., Chapman, W. E., Elmore, K., Fan, D., et al. (2024). Evidential deep learning: Enhancing predictive uncertainty estimation for Earth system science applications. *Artificial Intelligence for the Earth Systems*, 3(4). <https://doi.org/10.1175/AIES-D-23-0093.1>
- Umlauf, L., & Burchard, H. (2005). Second-order turbulence closure models for geophysical boundary layers. A review of recent work. *Continental Shelf Research*, 25(7), 795–827. <https://doi.org/10.1016/j.csr.2004.08.004>
- Williams, P. D. (2012). Climatic impacts of stochastic fluctuations in air–sea fluxes. *Geophysical Research Letters*, 39(10), 2012GL051813. <https://doi.org/10.1029/2012GL051813>
- Wu, J. (2025). jiarong-wu/mlflux: Code and model weights [Software]. *Zenodo*. <https://doi.org/10.5281/zenodo.17554656>
- Yu, L. (2019). Global air–sea fluxes of heat, fresh water, and momentum: Energy budget closure and unanswered questions. *Annual Review of Marine Science*, 11(1), 227–248. <https://doi.org/10.1146/annurev-marine-010816-060704>
- Zhou, S., Shi, R., Yu, H., Zhang, X., Dai, J., Huang, X., & Xu, F. (2024). A physical-informed neural network for improving air–sea turbulent heat flux parameterization. *Journal of Geophysical Research: Atmospheres*, 129(17), e2023JD040603. <https://doi.org/10.1029/2023JD040603>