

MIT Open Access Articles

Output-weighted and relative entropy loss functions for deep learning precursors of extreme events

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Rudy, Samuel H. and Sapsis, Themistoklis P. 2023. "Output-weighted and relative entropy loss functions for deep learning precursors of extreme events." *Physica D: Nonlinear Phenomena*, 443.

As Published: 10.1016/j.physd.2022.133570

Publisher: Elsevier BV

Persistent URL: <https://hdl.handle.net/1721.1/154215>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-ShareAlike



Output-weighted and relative entropy loss functions for deep learning precursors of extreme events

Samuel H. Rudy* and Themistoklis Sapsis

Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139

Abstract

Many scientific and engineering problems require accurate models of dynamical systems with rare and extreme events. Such problems present a challenging task for data-driven modelling, with many naive machine learning methods failing to predict or accurately quantify such events. One cause for this difficulty is that systems with extreme events, by definition, yield imbalanced datasets and that standard loss functions easily ignore rare events. That is, metrics for goodness of fit used to train models are not designed to ensure accuracy on rare events. This work seeks to improve the performance of regression models for extreme events by considering loss functions designed to highlight outliers. We propose a novel loss function, the adjusted output weighted loss, and extend the applicability of relative entropy based loss functions to systems with low dimensional output. The proposed functions are tested using several cases of dynamical systems exhibiting extreme events and shown to significantly improve accuracy in predictions of extreme events.

Keywords– Computational methods, machine learning, extreme events

1 Introduction

Accurate prediction and quantification of extreme events are critical tasks in many areas of science and engineering. Specific cases include the study of extreme weather patterns [16], turbulence [55], macroeconomic fluctuations [33], rouge waves [15], and many others [50, 20]. Recent research has developed tools for experimental design aimed towards uncertainty quantification in tail regions of systems with extreme events [39, 49, 5], prediction in the sense of classification of upcoming events in turbulent fluid flows [19, 6, 27], and regression problems for systems exhibiting extreme events [13, 43].

These recent works fall into the broader category of data-driven approaches to dynamical systems and fluid dynamics [7, 9]. Uses for such methods are motivated by cases in which physics based models fail due to intractable complexity, computational requirements, or insufficient measurements. In these cases machine learning and in particular deep learning offer a potential tool for improved predictive modelling. Successful applications of deep learning to problems in dynamical systems include flow reconstruction [38], physics informed neural networks [45], closure models [14, 26], sub-grid scale models [8, 48], climate modelling [30], operator inference [34], and embedding and lifting transformations [44, 35, 10].

Several recent works have explored the use of specific loss functions for training prediction models in the context of extreme events. Guth and Sapsis [27] develop the maximum adjusted area under the precision recall curve and show that it is effective in predicting extreme events in systems including the Kolmogorov flow and Majda-McLaughlin-Tabak model [36]. However, their proposed metric is not differentiable and not well approximated by small samples. Implementation

* Corresponding author (shrudy@mit.edu)

for high dimensional models such as neural networks would therefore be challenging. Doan et. al. [13] use a physics informed loss function improve the accuracy of echo-state networks for forecasting a Galerkin model of a turbulent flow with intermittent quasi-laminar states. While effective, this approach is constrained to problems where there is a known dynamic model for the quantity being predicted. Authors of [43] use a relative entropy based loss function to forecast the truncated Korteweg–de Vries equation, a simplified model of turbulent surface waves. This is shown to significantly improve performance, but requires a high dimensional target quantity. More recently, the use of various model architectures for predicting extreme events has been studied in [3].

In this work we seek to develop more broadly applicable loss functions and evaluate their performance on several challenging test problems. While the loss functions proposed in this work may be applied to arbitrary regression models, the included examples problems both employ neural networks. We assume the reader is familiar with common deep learning techniques including recurrent neural networks, stochastic optimization, and early stopping. The unfamiliar reader may find an excellent and free online reference in [24]. In particular, we make use of long-short-term-memory networks [29] for each of the test cases used in this work. The results could almost certainly be improved on via more carefully thought out network structures, training, and other user decisions [3]. However, such considerations are not the focus of this work which focuses solely on the effect of loss functions.

The paper is organized as follows; In Sec. 2 we outline the motivation for extreme event specific loss functions and develop methods including output weighted variations of the mean square error and a relative entropy method based on work in [43]. We also discuss error metrics for evaluating the proposed loss functions. In Sec. 3 we present results of the proposed loss functions applied to two test cases; Kolmogorov flow at Reynolds number $Re = 40$ and the flow around a square cylinder at $Re = 5000$. A discussion of the results, limitations, and outlook is presented in Sec. 4.

2 Methods

In this section we outline the proposed loss functions used for training neural networks to predict and quantify extreme events. These include two weighted variations of the mean square error as well as a relative entropy, also known as KL-divergence, based loss. We also describe methods for approximating the density function of the target variable y and the metrics we use to measure accuracy of the trained networks.

2.1 Problem statement

Let (X, Σ_x, μ) be a probability space with $X \subseteq \mathbb{R}^n$ and μ absolutely continuous with probability density function p_x . For some unknown function f , we have a dataset $\mathcal{D} = (\mathbf{X}, \mathbf{y})$ where $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$, $\mathbf{x}_i \sim p_x$, $\mathbf{y} = \{y_1, \dots, y_m\}$ with $y_i = f(\mathbf{x}_i)$, perhaps perturbed by measurement noise. We will make use of the fact that since $p(\mathcal{D}) = \prod_i p_x(\mathbf{x}_i)$ we have $\mathbb{E}_{\mathcal{D}}[\frac{1}{m} \sum g(\mathbf{x}_i)] = \mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]$ for any g such that the expectation is finite. That is, empirical averages over \mathcal{D} are unbiased approximations of integrals over p_x .

We are interested in parametric models \hat{f} approximating f which accurately predict and quantify outlier values of y and which accurately capture the induced probability density function through a measure transformation: $p_y(y) = d/dy \mu(f^{-1}(-\infty, y))$. To this end, both the structure of the model and training parameters have important effects. We seek to develop objective functions tailored for extreme event prediction that accurately capture extreme events, are differentiable, and may be approximated from finite datasets.

Throughout the remainder of the paper, we will use the terms true and false positive and neg-

ative to describe various results from a continuous regression model. In this context, we loosely define a true positive to be a prediction $\hat{f}(\mathbf{x})$ such that $p_y(f(\mathbf{x})) \sim p_y(\hat{f}(\mathbf{x})) \ll 1$. That is, when both the true and predicted values of y are rare. Likewise, a true negative is when $p_y(f(\mathbf{x})) \sim p_y(\hat{f}(\mathbf{x})) \sim \mathcal{O}(1)$, a false positive is when $p_y(f(\mathbf{x})) \sim \mathcal{O}(1)$ and $p_y(\hat{f}(\mathbf{x})) \ll 1$, and a false negative is when $p_y(f(\mathbf{x})) \ll 1$ and $p_y(\hat{f}(\mathbf{x})) \sim \mathcal{O}(1)$. In some places, these definitions will be made rigorous by applying thresholds to $p_y(f(\mathbf{x}))$ and $p_y(\hat{f}(\mathbf{x}))$.

2.2 Output-weighted variations of mean square error

The mean square error is the most common loss function used for training regression models with real valued outputs. It is given by,

$$L_{MSE}(\hat{f}) = \mathbb{E}_{\mathbf{x}} \left[e_{\hat{f}}(\mathbf{x})^2 \right] = \int_{\mathcal{X}} e_{\hat{f}}(\mathbf{x})^2 p_{\mathbf{x}}(\mathbf{x}) d\mathbf{x} = \mathbb{E}_{\mathcal{D}} \left[\frac{1}{m} \sum_{\mathcal{D}} e_{\hat{f}}(\mathbf{x}_i)^2 \right], \quad (1)$$

where $e_{\hat{f}} = (f(\mathbf{x}) - \hat{f}(\mathbf{x}))$. Squaring the error makes L_{MSE} more sensitive to true outliers than the mean absolute error, but if outlier values make up a small fraction of the total dataset then L_{MSE} may still be small while missing large y . Moreover, if some rare values of y are not separated from the core of p_y by substantial distance then they may be missed with little added error. To better see this, consider the case where μ admits a disintegration over the induced measure on y . Then we can express L_{MSE} as an integral over Y of the regular conditional expectation of the square error. That is,

$$L_{MSE}(\hat{f}) = \mathbb{E}_y \left[\mathbb{E}_{\mathbf{x}} \left[e_{\hat{f}}(\mathbf{x})^2 \mid f(\mathbf{x}) = y \right] \right] = \int_Y p_y(y) \mathbb{E}_{\mathbf{x}} \left[e_{\hat{f}}(\mathbf{x})^2 \mid f(\mathbf{x}) = y \right] dy, \quad (2)$$

where the conditional expectation is defined using the disintegration of $p_{\mathbf{x}}$ over p_y [11]. For rare events, the value of $p_y(y)$ is small, allowing large error in the prediction of such events without significantly affecting L_{MSE} .

The insensitivity of L_{MSE} to rare events may be mitigated via introducing a weighting function. Specifically, consider the case where the square error is weighted according to $p_y(f(\mathbf{x}))^{-1}$. The resulting function is given by,

$$\begin{aligned} L_{OW}(\hat{f}) &= \mathbb{E}_{\mathbf{x}} \left[p_y(f(\mathbf{x}))^{-1} e_{\hat{f}}(\mathbf{x})^2 \right] \\ &= \int_Y p_y(y) \mathbb{E}_{\mathbf{x}} \left[p_y(f(\mathbf{x}))^{-1} e_{\hat{f}}(\mathbf{x})^2 \mid f(\mathbf{x}) = y \right] dy \\ &= \int_Y \mathbb{E}_{\mathbf{x}} \left[e_{\hat{f}}(\mathbf{x})^2 \mid f(\mathbf{x}) = y \right] dy. \end{aligned} \quad (3)$$

Note that the expression $p_y(f(\mathbf{x}))|_{f(\mathbf{x})=y}$ is simply $p_y(y)$ and therefore cancels the $p_y(y)$ term outside the conditional expectation. We call the expression given by Eq. (3) the output-weighted loss, L_{OW} since the square error is weighed by the inverse of the likelihood of the true output $f(\mathbf{x})$. Expressions with similar form have been used for sequential sampling strategies for rare events [49, 5]. However, these works were focused on experimental design and used least squares or Gaussian process regression to model f . Equation (3) is also related to oversampling techniques commonly used for classification problems with imbalanced data [28].

As a cost function, Eq. (3) has some potential drawbacks. In particular, the weight given to each error is proportional only to the inverse likelihood of the true value, $p_y(f(\mathbf{x}))$, and is independent from $p_y(\hat{f}(\mathbf{x}))$. Thus, error accumulated on false negatives is penalized far more than that made on




Figure 1: Visualization of weights used by L_{OW} and L_{AOW} . Left most figure shows hypothetical boundaries for true/false positive/negative regions. Pseudocolor plots show output weight, adjusted output weight, and difference. Disagreement is primarily within the false positive region.

false positives. As we will show in Sec. 3, minimizing Eq. (3) often yields models that over-predict rare events.

The problem of false positives may be addressed with the inclusion of a second term weighing the square error. The weight function $1/p_y(f(\mathbf{x}))$ in Eq. (3) amplifies any error realized on examples with rare y . This included true positive and false negative predictions. To distinguish between true negative and false positive predictions, we require a weight function that depends on $p_y(\hat{f}(\mathbf{x}))$. The ratio $p_y(f(\mathbf{x}))/p_y(\hat{f}(\mathbf{x}))$ is large only in the case where the likelihood of the predicted output is lower than that of the true output. Thus, the expression given by,

$$L_{FP}(\hat{f}) = \mathbb{E}_{\mathbf{x}} \left[\frac{p_y(f(\mathbf{x}))}{p_y(\hat{f}(\mathbf{x}))} e_{\hat{f}}(\mathbf{x})^2 \right], \quad (4)$$

is a measure of the error made on false positive predictions. Summing Eq. (3) with Eq. (4) gives,

$$\begin{aligned} L_{AOW}(\hat{f}) &= L_{OW}(\hat{f}) + L_{FP}(\hat{f}) \\ &= \mathbb{E}_{\mathbf{x}} \left[\left(\frac{1}{p_y(f(\mathbf{x}))} + \frac{p_y(f(\mathbf{x}))}{p_y(\hat{f}(\mathbf{x}))} \right) e_{\hat{f}}(\mathbf{x})^2 \right] \end{aligned} \quad (5)$$

which we will call the *adjusted output weighted loss*, or L_{AOW} . For true positives, when both $p_y(f(\mathbf{x}))$ and $p_y(\hat{f}(\mathbf{x}))$ are small, the second term in the parentheses is $\mathcal{O}(1) \ll p_y(f(\mathbf{x}))^{-1}$. The integrand therefore agrees with that of Eq. (3). For true negatives and false negatives, the integrand is also of similar magnitude to that in Eq. (3). However, for false positives, the second term in the parenthesis is large. The expression therefore penalizes errors made as false positive predictions to a greater degree than Eq. (3).

The added penalization of false positives comes at the cost of increased complexity. Equation (3) is a weighted least squares problem and does not add any computational complexity to Eq. (1) beyond pre-computing $\{p_y(y_i)\}_{i=1}^m$. In particular, Eq. (3) has closed form solution for linear problems. This is not the case for Eq. (5), where iterative optimization is required. For \hat{f} depending nonlinearly on parameters, such as neural networks, the difference is less important, as iterative methods must be used in either case.

Differences between the weighting functions in the integrands of Eq. (3) and (5) are visualized in Fig. 1. We plot values of the each term using hypothetical values of $p_y(f(\mathbf{x}))$ and $p_y(\hat{f}(\mathbf{x}))$. The left most plot illustrates a rough partitioning of the domain into true/false positives and negatives. Note that many expressions of the likelihoods of the true and predicted y could be constructed to

find novel loss functions. The authors suspect that this would be a fruitful research direction, but it is not within the scope of this work.

2.3 Approximating p_y

Computation of the two weighted variations of the mean square error in the previous section requires evaluating p_y at points y_i in the training set. For Eq. (3), this may be done offline using any off-the-shelf density estimation technique. Evaluating (5) further requires $p_y \circ \hat{f}$, which changes during training and must be evaluated rapidly. Gradient based optimization also requires evaluating $p'_y \circ \hat{f}$. Thus, minimizing Eq. (5) requires a low-computational-cost differentiable approximation of p_y . This rules out non-parametric density estimates such as kernel density estimation (due to slowness), histograms (due to lack of differentiability), and k-nearest-neighbors (slow and not differentiable). We instead use a combined method of nonparametric estimation of $\log(p_y(y_i))$ which we then fit using a Gaussian process with few collocation points, allowing for rapid evaluation of p_y and its derivative.

Estimates of $p_y(y_i)$ are initially made using a histogram of $\{y_i\}_{i=1}^m$ with $n_b = 100$ evenly sized bins. Defining $B_i = [b_i, b_{i+1})$ for $i = 1, \dots, 99$ and $B_{100} = [b_{100}, b_{101}]$ where b_i are evenly spaced values between $\min(y_i)$ and $\max(y_i)$ we have for $y \in B_i$,

$$\log \hat{p}(c_i) \approx \log \left(\frac{|\{j : y_j \in B_i\}|}{ml} \right) \quad (6)$$

where $l = b_{i+1} - b_i$ and $c_i = (b_i + b_{i+1})/2$.

Following estimation of the log densities via histogram, the values at the center of each bin are fit to a Gaussian process [47]. We use the Matern-5/2 kernel with additional white heteroscedastic noise. The mean function of the Gaussian process is set to be the log of machine epsilon, enforcing that the approximation goes to zero away from the sampled data. That is,

$$\log \hat{p}_y(y) \sim \mathcal{GP} \left(\mu, k(y, y') + \sigma_w^2(y) \delta_{y, y'} \right), \quad (7)$$

where $\mu = \log(10^{-16})$, $\sigma_w(y)$ is the white noise term and $\delta_{y, y'}$ is the Kronecker delta function. Fitting the logarithm of the probability rather than the probability helps to get accurate estimates in the tails of p_y . We obtain estimates of the density at sample points y using the conditional mean of the Gaussian process;

$$\begin{aligned} \hat{p}_y(y) &= \exp \left(k(y, \mathbf{c}) \left(k(\mathbf{c}, \mathbf{c}) + \sigma_w^2(\mathbf{c}) \mathbf{I} \right)^{-1} \log \hat{p}_y(\mathbf{c}) + \mu \right) \\ &= \exp \left(k(y, \mathbf{c}) \boldsymbol{\alpha} + \log(10^{-16}) \right). \end{aligned} \quad (8)$$

where $c_i = (b_i + b_{i+1})/2$ and vector $\boldsymbol{\alpha}$ is pre-computed and stored. The expected gradient of the exponent is given by simply differentiating the kernel [37]. This allows for queries $\hat{p}_y(y)$ and its gradient with the simple computation of $k(y, \mathbf{c})$ and $k'(y, \mathbf{c})$. In practice we found training to be more stable when a floor was set for the value of $p_y(y)$. Examples in this work all used an effective density equal to $\hat{p}_y(y) + 10^{-5}$ which has negligible effect of most events.

We note that the same process could be easily implemented for y with dimension greater than one but still low. Initial estimates of p_y at training points of the Gaussian process could be taken with any standard non-parametric density estimation [54] and subsequently fit to a GP. However, for higher dimensional y , both density approximation and Gaussian process interpolation become non-trivial. It is possible that in these cases the density p_y could be substituted for that of a relevant observable $g(y)$, but such work is beyond the scope of the present manuscript.

2.4 A relative entropy based loss function

The use of relative entropy as a loss function for neural networks was explored in [43]. Qi and Majda used the relative entropy (i.e. KL-divergence) between truth and prediction, after applying the soft-max function. Specifically, for $y \in \mathbb{R}^s$, [43] uses a loss function defined for a single datapoint by,

$$\begin{aligned} L_{QM}(\mathbf{x}) &= KL\left(\sigma(f(\mathbf{x})) \parallel \sigma(\hat{f}(\mathbf{x}))\right) + \alpha KL\left(\sigma(-f(\mathbf{x})) \parallel \sigma(-\hat{f}(\mathbf{x}))\right) \\ &= \sum \left(\sigma(f(\mathbf{x})) \log \left(\frac{\sigma(f(\mathbf{x}))}{\sigma(\hat{f}(\mathbf{x}))} \right) + \alpha \sigma(-f(\mathbf{x})) \log \left(\frac{\sigma(-f(\mathbf{x}))}{\sigma(-\hat{f}(\mathbf{x}))} \right) \right) \end{aligned} \quad (9)$$

where log is taken elementwise, the sum is over s dimensions of the vector enclosed in the parenthesis, and the soft-max function, σ , is defined by,

$$\sigma(y) = \frac{\exp(y)}{\sum \exp(y)} \quad (10)$$

The use of σ weights outputs by the exponent of their magnitude, thus ensuring the loss focuses on accurate learning of large magnitude features. This approach is extremely effective in [43], where y the solution to a PDE and thus high dimensional. However, it may not be applied directly in the case of scalar output y . This is because for any $y \in \mathbb{R}^1$, $\sigma(y) = 1$. Moreover, Eq. (9) is only able to weight extreme values within one output sample, rather than comparing multiple y . However, it is possible to derive similar loss functions where a soft-max like operator is applied to multiple samples, rather than the indices within a single sample.

Let us assume that f and \hat{f} are such that $\mathbb{E}_x[\exp(f(\mathbf{x}))]$ and $\mathbb{E}_x[\exp(\hat{f}(\mathbf{x}))]$ are finite and non-zero. Note that this is a mild assumption that holds on a set containing $L^\infty(X)$. Then we can define an operator $G(f)$ by,

$$\begin{aligned} G(f)(\mathbf{x}) &= e^{f(\mathbf{x})} p_x(\mathbf{x}) \left(\int_X e^{f(\mathbf{x})} p_x(\mathbf{x}) d\mathbf{x} \right)^{-1} \\ &= e^{f(\mathbf{x})} p_x(\mathbf{x}) \mathbb{E}_x \left[e^{f(\mathbf{x})} \right]^{-1}. \end{aligned} \quad (11)$$

This operator acts as a continuous analog of the soft-max function. Functions in the range of G are probability density functions on X where the value of $G(f)(\mathbf{x})$ is proportional to the sample density $p_x(\mathbf{x})$ and the exponent of $f(\mathbf{x})$. $G(f)$ therefore has mass concentrated on those values of \mathbf{x} whose likelihood under p_x is not vanishing and whose image under f is large, or in other words, extreme. Note that in the case where f is linear the operator defined in Eq (11) is known as exponential tilting [17], which has previously been used in importance sampling [51]. We define the relative entropy loss as,

$$L_{RE}(\hat{f}) = KL\left(G(f) \parallel G(\hat{f})\right). \quad (12)$$

Note that compared to Eq. (9) used in [43], the normalization of $G(f)$ used in Eq. (12) is taken across the input space X rather than dimensions of y . This allows for exponential weighting of outputs by their magnitude even in the case of scalar y . We are interested in minimizing Eq. (12) with respect to \hat{f} . Expanding Eq. (12) and ignoring terms whose value does not depend on \hat{f} we find,

$$\begin{aligned} L_{RE}(\hat{f}) &= \mathbb{E}_x \left[\frac{e^{f(\mathbf{x})}}{\mathbb{E}_x \left[e^{f(\mathbf{x})} \right]} \log \left(\frac{e^{f(\mathbf{x})} p_x(\mathbf{x}) \mathbb{E}_x \left[e^{\hat{f}(\mathbf{x})} \right]}{e^{\hat{f}(\mathbf{x})} p_x(\mathbf{x}) \mathbb{E}_x \left[e^{f(\mathbf{x})} \right]} \right) \right] \\ &\propto \mathbb{E}_x \left[e^{f(\mathbf{x})} \log \left(e^{f(\mathbf{x}) - \hat{f}(\mathbf{x})} \mathbb{E}_x \left[e^{\hat{f}(\mathbf{x})} \right] \right) \right] \\ &= \mathbb{E}_x \left[e^{f(\mathbf{x})} (f(\mathbf{x}) - \hat{f}(\mathbf{x})) + e^{f(\mathbf{x})} \log \left(\mathbb{E}_x \left[e^{\hat{f}(\mathbf{x})} \right] \right) \right]. \end{aligned} \quad (13)$$

Individual expectations in the above expression may be estimated with sums over the dataset. However, the term inside the log is problematic. By Jensen's inequality,

$$\log\left(\mathbb{E}_{\mathbf{x}}\left[e^{\hat{f}(\mathbf{x})}\right]\right) = \log\left(\mathbb{E}_{\mathcal{D}}\left[\frac{1}{m}\sum_{i=1}^m e^{\hat{f}(\mathbf{x}_i)}\right]\right) \geq \mathbb{E}_{\mathcal{D}}\left[\log\left(\frac{1}{m}\sum_{i=1}^m e^{\hat{f}(\mathbf{x}_i)}\right)\right] \quad (14)$$

where the second and third expectations are over the random samples \mathcal{D} . Thus, the expected log of the empirical average of $\exp(\hat{f}(\mathbf{x}))$ is an underestimate. We therefore find an upper bound for Eq. (14) that may be accurately approximated and minimized. Note that for any α , log is bounded above by its first order Taylor expansion about α . Therefore,

$$\begin{aligned} \log\left(\mathbb{E}_{\mathbf{x}}\left[e^{\hat{f}(\mathbf{x})}\right]\right) &= \log\left(\alpha + \left(\mathbb{E}_{\mathbf{x}}\left[e^{\hat{f}(\mathbf{x})}\right] - \alpha\right)\right) \\ &\leq \log(\alpha) + \frac{\mathbb{E}_{\mathbf{x}}\left[e^{\hat{f}(\mathbf{x})}\right] - \alpha}{\alpha}. \end{aligned} \quad (15)$$

The error in the Taylor expansion and thus tightness of the bound is on the order of $(\mathbb{E}_{\mathbf{x}} \exp(\hat{f}(\mathbf{x})) - \alpha)/\alpha^2$. We therefore want to pick some α as close to $\mathbb{E}_{\mathbf{x}} \exp(\hat{f}(\mathbf{x}))$ as possible. Consider $\alpha = \mathbb{E}_{\mathbf{x}} \exp(f(\mathbf{x}))$ which for $f \approx \hat{f}$ we assume will be close. In this case the upper bound for L_{RE} simplifies dramatically to,

$$\begin{aligned} L_{RE}(\hat{f}) &\leq \mathbb{E}_{\mathbf{x}} \left[e^{f(\mathbf{x})}(f(\mathbf{x}) - \hat{f}(\mathbf{x})) + e^{f(\mathbf{x})} \left(\log\left(\mathbb{E}_{\mathbf{x}}\left[e^{f(\mathbf{x})}\right]\right) + \frac{\mathbb{E}_{\mathbf{x}}\left[e^{\hat{f}(\mathbf{x})}\right] - \mathbb{E}_{\mathbf{x}}\left[e^{f(\mathbf{x})}\right]}{\mathbb{E}_{\mathbf{x}}\left[e^{f(\mathbf{x})}\right]} \right) \right] \\ &= \mathbb{E}_{\mathbf{x}} \left[e^{\hat{f}(\mathbf{x})} - e^{f(\mathbf{x})} \hat{f}(\mathbf{x}) \right], \end{aligned} \quad (16)$$

where we have ignored terms that do not depend on \hat{f} . In particular, the log expectation term, $\log(\alpha)$, has been removed since it does not depend on parameters. The remaining terms are easily approximated from dataset \mathcal{D} by,

$$L_{RE}(\hat{f}) = \mathbb{E}_{\mathbf{x}} \left[e^{\hat{f}(\mathbf{x})} - e^{f(\mathbf{x})} \hat{f}(\mathbf{x}) \right] = \mathbb{E}_{\mathcal{D}} \left[\frac{1}{m} \sum_{i=1}^m \left(e^{\hat{f}(\mathbf{x}_i)} - e^{f(\mathbf{x}_i)} \hat{f}(\mathbf{x}_i) \right) \right]. \quad (17)$$

Following [43] we note that the relative entropy loss only focuses the error on large positive values of $f(\mathbf{x})$ and introduce the generalization,

$$L_{RE,\lambda}(\hat{f}) = L_{RE}(\hat{f}) + \lambda L_{RE}^{(-)}(\hat{f}), \quad (18)$$

where $L_{RE}^{(-)}(\hat{f})$ is defined by replacing f and \hat{f} in Eq. (17) with $-f$ and $-\hat{f}$ respectively. The value of λ is a tuning parameter that can be set according to the skew of the dataset. Data considered in this work contains extremes that skew positive. We therefore set $\lambda = 0.1$, so that the loss function focuses on predicting positive outliers.

2.5 Performance measures for regression with extreme events

Goodness of fit in the context of regression for extreme events is non-trivial to define. In Sec. 2.1 we outlined three criteria for a "good" predictor. That is, we seek models that accurately predict extreme events (in the sense that they may be used as a classifier), quantify extreme events, and yield accurate densities in the tails of p_y . In this section we present metrics for quantifying each of the three criteria and discuss potential shortcomings of our approach towards error analysis.

2.5.1 Accurate quantification of tails of $p_y(y)$

Finally, we seek models such that the push-forward density under the learned model is similar to the true density p_y . We are particularly interested in loss functions yielding densities that match the tails of the true density. Following [49, 5] use the difference between logarithms of the two density functions over the intersection of their support. This is normalized by the size of the intersection to penalize distributions that only intersect on a small domain. The metric is given by;

$$\mathbb{D}(p_y, \hat{p}_y) = \frac{1}{|\Omega(p_y, \hat{p}_y)|^2} \int_{\Omega(p_y, \hat{p}_y)} |\log(p_y(y)) - \log(\hat{p}_y(y))| dy, \quad (19)$$

where

$$\Omega(p_y, \hat{p}_y) \approx \text{supp}(p_y) \cap \text{supp}(\hat{p}_y). \quad (20)$$

and \hat{p}_y is the density under the learned model. Since p_y and \hat{p}_y are approximated from data, their support is unknown and behavior in low density regions extremely challenging to quantify. Unfortunately, \mathbb{D} depends strongly on both of these quantities. We therefore approximate the support of each distribution as the interval covering the observed range of values. This is an underestimate of the true width, but allows for a consistent method of computing \mathbb{D} .

2.5.2 Accurate quantification of extreme events

Accurate quantification is indicated by models achieving low error on predictions from $\mathbf{x} \in X$ such that $f(\mathbf{x})$ is rare. We quantify this with the expected mean square error over the set of inputs corresponding to rare events. For example, consider the mean square error restricted to the set of events with $p_y(y) < \epsilon$;

$$MSE_\epsilon = \mathbb{E}_{\mathbf{x}} [e(\mathbf{x})^2 | p_y(f(\mathbf{x})) \leq \epsilon], \quad (21)$$

for values of $\epsilon > 0$. Models that accurately quantify rare events should have low MSE_ϵ for $\epsilon \ll 1$. A more informative metric would condition on $p_y(f(\mathbf{x})) = \epsilon$, but computation of such a quantity requires a parametric model or smoother. We therefore do not include it in this work.

2.5.3 Accurate prediction of extreme events

Models trained for regression may be used along with some threshold value to function as classifiers. For a variety of extreme event rates, we track two metrics of classifier accuracy. Following [27], we use the extreme event rate dependent area under the precision recall curve given by,

$$\alpha(\omega; y, \hat{y}) = \int_{\mathbb{R}} s(1_{y \geq a}, 1_{\hat{y} \geq b}) \left| \frac{\partial}{\partial b} r(1_{y \geq a}, 1_{\hat{y} \geq b}) \right| db \quad (22)$$

where s and r are the precision and recall and $a = F_y^{-1}(1 - \omega)$ where F_y^{-1} is the quantile function for y . This quantity has in fact been used to train models with lower dimensional parameter spaces. However, the use of distinct thresholds for y and \hat{y} may be undesirable. Note for example that for any strictly increasing function g we have $\alpha(\omega; y, \hat{y}) = \alpha(\omega; y, g(\hat{y}))$. We therefore also consider a metric that uses the same threshold for both y and \hat{y} . The extreme event rate dependent F_1 score is given by,

$$F_1(\omega; y, \hat{y}) = F_1(1_{y \geq a}, 1_{\hat{y} \geq a}) \quad \text{where} \quad a = F_y^{-1}(1 - \omega) \quad (23)$$

where the F_1 -score is given by the harmonic mean of precision and recall.

In many practical settings, the consequences of false positive or negative predictions will be disproportionate. In these cases, the metrics given by Eq. (22) and (23) will not reflect the utility of the learned model and more setting specific metrics should be considered.

3 Results

In this section we present the results of using each of the loss functions discussed in Sec. 2 to two challenging supervised learning problems resulting from fluid dynamic systems with extreme events; the Kolmogorov flow at Reynolds number 40, and the flow around a square cylinder at Reynolds number 5000. Numerical simulations of the incompressible Navier-Stokes equations in each case were performed using the spectral element method implemented in Nek5000 [41, 42], an open source Fortran code for incompressible fluids. Details on numerical simulation and data preparation are given in Appendix A. In both cases, target data y is centered and normalized to have zero mean and unit variance.

In each case we use LSTM networks implemented in tensorflow [1] and trained using Adam [31] with early stopping to avoid over-fitting. Datasets are split into training (50%), validation (10%) and testing (40%). Each partition is formed from a contiguous set of samples so that phenomena observed in each are distinct. Further details on training, as well as network structure are given in Appendix B. We report the metrics outlined in Sec. 2.5 evaluated on the portion of the data reserved for testing.

For each result we present data from twenty randomly initialized and trained networks. Plots show mean value of across all trials as a solid line and shade region between 10th and 90th percentiles, thus excluding the two highest and two lowest values.

3.1 Kolmogorov Flow

We first consider two dimensional Kolmogorov flow at $Re = 40$. Dynamics follow the incompressible Navier-Stokes equations with a sinusoidal forcing term and periodic boundaries. Specifically,

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} + \sin(k_f y) \mathbf{e}_1, \quad \nabla \cdot \mathbf{u} = 0, \quad (24)$$

where $\mathbf{e}_1 = (1, 0)$, $k_f = 4$, and where we have boundary conditions $\mathbf{u}(2\pi, y, t) = \mathbf{u}(0, y, t)$ and $\mathbf{u}(x, 2\pi, t) = \mathbf{u}(x, 0, t)$. We are interested in this flow due to the behavior exhibited by the energy dissipation rate, given by a re-scaling of the enstrophy,

$$D(t) = \frac{\nu}{|\Omega|} \int_{\Omega} |\nabla \mathbf{u}(\mathbf{z}, t)|^2 dz, \quad (25)$$

where $\nu = Re^{-1}$ is the viscosity and $\Omega = [0, 2\pi)^2$ is the computational domain. Solutions to Eq. (24) with the prescribed boundary conditions are known to exhibit large bursts in energy input and dissipation rate resulting from intermittent alignment of the velocity field with the external forcing [18]. An single snapshot of the velocity field in the x-direction, as well as the time series for the energy dissipation rate and its density are shown in Fig. 2.




Figure 2: Kolmogorov flow. Left: snapshot of x-velocity. Center: Time series for energy dissipation. Right: Empirical density function of energy dissipation for $t \in [0, 40,000]$.




Figure 3: Mean predictions for subset of test data for the Kolmogorov dataset at lead time $\tau = 6$. Blue time series is true $D(t)$ and red is mean prediction across 20 trained neural networks.

The intermittent bursts apparent in Fig. 2 have been the subject of several previous works seeking to predict their occurrence in advance [18, 27, 3]. In particular, Farazmand and Sapsis [19] showed that the intermittent behavior could be explained through triadic interactions between Fourier modes. Here we consider the same problem, but seek to predict bursts in the energy dissipation using neural networks. The target (y) quantity is the energy dissipation rate $D(t)$, normalized to have zero mean and unit variance. We use the time varying magnitudes of the three Fourier modes identified in [19] as inputs; $\mathbf{x}(t) = (a_{0,k_f}(t), a_{1,0}(t), a_{1,k_f}(t))$. In order to ensure accurate statistical descriptions of errors, a simulation of the Kolmogorov flow run for $4 \cdot 10^4$ time units.




Figure 4: Density for the Kolmogorov flow energy dissipation using true data (blue) and 20 realizations of neural networks trained with each loss function (red). Reported \mathbb{D} is mean of Eq. (19).

We train neural network predictors of $D(t + \tau)$ from $\{\mathbf{x}(s) : s \leq t\}$ for several lead times τ . LSTM networks [29] are used to allow for historical data to assist in prediction. Further details on the neural network structure and training procedure may be found in Appendix B. For each

network configuration and lead time, we train twenty networks from randomly initialized weights. A subset of the test data results is shown for a lead time of $\tau = 6$ in Fig. 3. The true value of $D(t)$ is shown in blue and the mean prediction across twenty networks for each loss function is shown in red. A cursory inspection reveals that L_{MSE} appears to underestimate the large fluctuations, though it does appear to predict their locations. Each of the three other loss functions appear to quantify the peaks more accurately, with L_{OW} having the most pronounced false positives.

Estimates of the densities p_y using test set values of $y_i = f(\mathbf{x}_i)$ and predictions $\hat{y}_i = \hat{f}(\mathbf{x}_i)$ with \hat{f} trained using each of the loss functions considered in this work are shown in Fig. 4. The normalized difference between logs metric \mathbb{D} is also shown alongside each loss function and lead time. At each lead time, the induced density using models trained with L_{MSE} has the highest fidelity in the core of the pdf, but these distributions dramatically underestimate the density in the high dissipation tail region. At a low lead time of $\tau = 2$, densities from models trained using L_{RE} , L_{OW} , and L_{AOW} perform approximately equally. For higher τ , L_{RE} and L_{AOW} outperform L_{OW} , with the latter tending to overestimate the density positive outliers. Excepting $\tau = 10.0$, the minimal value of Eq. (19) is obtained by L_{AOW} . At $\tau = 10$ the best numerical result is obtained using L_{MSE} . However, the estimated density function reveals significant underestimation of tail events.




Figure 5: Comparison of MSE_ϵ for models trained using L_{MSE} (blue) and extreme event specific loss functions (red) for the Kolmogorov dataset. Note that high $-\log(\epsilon)$ corresponds to rare events.




Figure 6: Comparison of $\alpha(\omega)$ for models trained using L_{MSE} (blue) and extreme event specific loss functions (red) for the Kolmogorov dataset.

The rare event error given by Eq. (21) is shown in Fig. 5. Each subplot includes error statistics

for twenty networks trained with L_{MSE} in blue and networks trained using each of the specialized loss functions in red. In all cases, L_{MSE} has lower error when ϵ is small, which is unsurprising since this is the objective L_{MSE} minimizes. The difference is more pronounced when compared to L_{OW} or when τ is small. For rare events, each of the specialized loss functions has lower error.




Figure 7: Comparison of $F_1(\omega)$ for models trained using L_{MSE} (blue) and extreme event specific loss functions (red) for the Kolmogorov dataset. Lines indicates mean and shaded region is bounded by 10th and 90th percentiles.

The performance of the trained networks as classifiers for rare events is shown in Figs. 6 and 7. Fig. 6 shows that area under the precision recall curve as a function of extreme event rate. At low lead times and extreme event rates, L_{OW} and L_{AOW} show some improvement over MSE, but all methods are approximately equivalent in other cases. Note however, that $\alpha(\omega)$ allows for alternative thresholds to be applied to y and \hat{y} . Fig. 7 shows very clearly that across all cases the extreme event specific loss functions yield models that are more accurately able to distinguish extreme events when the same threshold is applied.

3.2 Flow around a square cylinder

Our second example considers the flow around a square cylinder at Reynolds number $Re = 5000$. The square has unit side-length and is positioned in a stream having unit inlet velocity and viscosity $\nu = 2 \cdot 10^{-4}$. The flow is characterized by extremely chaotic vortex shedding in the wake of the cylinder, shown in the left panel of Fig. 8 and large deviations in the forcing applied to the cylinder by the fluid. The drag coefficient is a non-dimensional quantity given by,

$$C_d(t) = \frac{2}{\rho u_\infty^2 C} \oint_{\partial S} (\boldsymbol{\tau}(t) - p(t)\mathbf{n}) \cdot \mathbf{e}_x ds \quad (26)$$

where $\boldsymbol{\tau}$, p , and \mathbf{n} are the skin shear stress, pressure, and normal vector, and ∂S is the boundary of the square cylinder. ρ , u_∞ , and C all have numerical value of 1 and are included only for their dimension. The time series of $C_d(t)$ for a simulation of length $T = 2 \cdot 10^4$ and empirical density function are shown in the center and right panels of Fig. 8.




Figure 8: Flow around a square cylinder at $Re = 5000$. Left: Snapshot of vorticity close to cylinder. Center: Time series for drag coefficient. Right: Empirical density function of drag coefficient.

We consider the problem of predicting $C_d(t + \tau)$, centered and normalized to unit variance, from 40 evenly spaced measurements of pressure on the surface of the cylinder, $P(t) \in \mathbb{R}^{40}$. The contributions of τ and $p(t)$ to C_d are called the skin friction drag and pressure drag, respectively. In this example, pressure drag is several orders of magnitude larger than skin friction drag. Therefore, assuming a sufficiently dense distribution of pressure measurements prediction of C_d at zero lead time is proportional to the difference of average pressure on the front and back of the square. However, for $\tau > 0$ the problem rapidly becomes challenging.

We test of of the loss functions proposed in Sec. 2 using an LSTM network described in further detail in Appendix B. Time series of the true drag coefficient as well as the mean prediction from twenty neural networks trained using each of the loss functions are shown in Fig. 9 for a lead time of $\tau = 1$. It is immediately clear that predictions made by networks trained with different loss functions exhibit substantially different behavior. Consistent with the Kolmogorov flow data, L_{MSE} substantially underestimates fluctuations, while L_{OW} has many false positives.




Figure 9: True centered and normalized drag coefficient (blue) and mean time series predictions (red) for test set data from each neural network for the drag on a square cylinder. Results are for a lead time of $\tau = 1$.

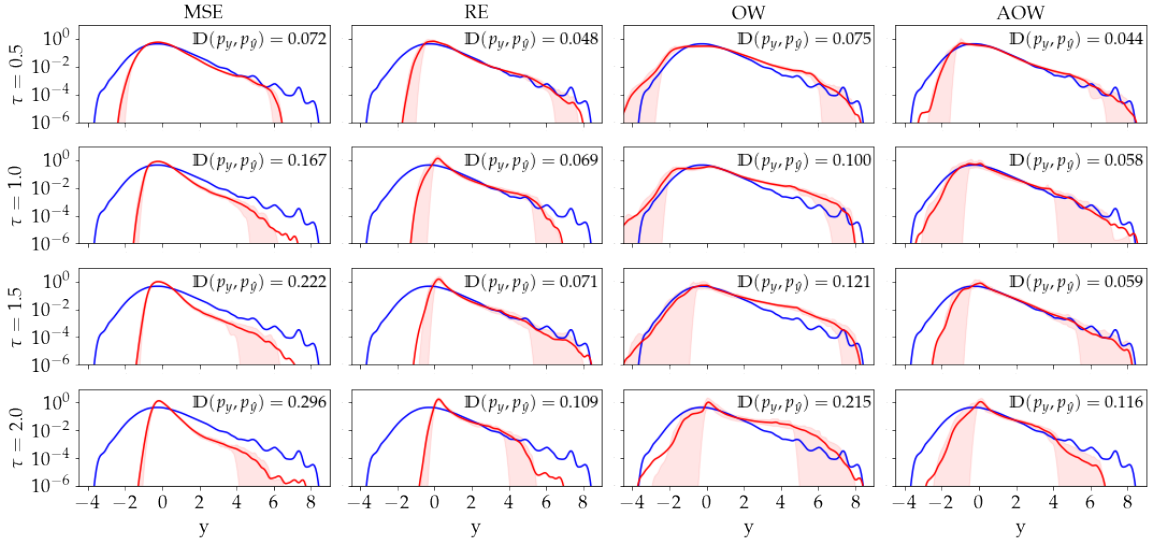


Figure 10: Density functions for the drag on a square cylinder using true data (blue) and 20 realizations of neural networks trained with each loss function (red).

Estimates of the density of y and of \hat{y} formed from test set data for each loss function and several lead times are shown in Fig. 10. As expected, networks trained with L_{MSE} underestimate the likelihood of events in the tails of the distribution of y while L_{OW} overestimates these same events. L_{RE} and L_{AOW} generally both capture the correct tail behavior for extreme events skewing positive, with L_{AOW} achieving lower \mathbb{D} . We note that while L_{RE} performs poorly for quantifying p_y when $y < 0$, this may be a consequence of our choice of parameter λ , which weighed positive values ten times as much as negative. In each case, L_{RE} and L_{AOW} notably outperform L_{OW} which in turn outperforms MSE with regards to mean \mathbb{D} .

Plots of MSE_ϵ are shown in Fig. 11. As expected loss functions tailored to extreme events are better able to quantify events occurring with a low probability (high $-\log(\epsilon)$), while underperforming L_{MSE} for more frequent events. This is consistent with expectations and supports the use of specialized loss functions for quantifying rare events.




Figure 11: Comparison of MSE_ϵ for models trained using L_{MSE} (blue) and extreme event specific loss functions (red) for the square dataset. Note, high $-\log(\epsilon)$ corresponds to rare events.




Figure 12: Comparison of $\alpha(\omega)$ for models trained using L_{MSE} (blue) and extreme event specific loss functions (red) for the square dataset.




Figure 13: Comparison of $F_1(\omega)$ for models trained using L_{MSE} (blue) and extreme event specific loss functions (red) for the square dataset.

Classification metrics for the square data are shown in Figs. 12 and 13. Figure 12 provides evidence that challenges the utility of the extreme event specific loss functions. At a lead time of $\tau = 1$, L_{RE} slightly outperforms L_{MSE} in the low extreme event rate region. However, L_{MSE} has the highest $\alpha(\omega)$ in all other cases. This suggests that while networks trained L_{MSE} do a poor job at quantifying extreme events for this example, they are in fact able to do a better job than others at separating extreme from quiescent events given an appropriate pair of thresholds. However, this is not true when the same threshold is applied to each of y and \hat{y} . Figure 13 shows that each of the extreme event specific loss functions outperforms L_{MSE} when the same threshold is used.

4 Discussion

In this work we have developed and evaluated several candidate loss functions for use in regression problems seeking to accurately quantify and predict extreme events. We have taken care to describe the motivation for each of the loss functions as an approximation of a continuous functional, allowing for greater intuition and providing a framework upon which further improvements may be made. The test cases used to evaluate the loss functions include the Kolmogorov flow, which has been studied extensively as a canonical example of a turbulent fluid flow exhibiting extreme events and the flow around a square cylinder at $Re = 5000$. This latter example has, to the best of our knowledge, not been studied in the context of extreme events. In each case, the proposed loss functions yield models with significantly improved error in the tails of the output, more accurate

estimates of the tail likelihoods, and improved classification when an equal threshold is applied to prediction and ground truth.

We have also provided evidence, however, that prediction and quantification of extreme events are distinct problems. Figure 12 makes it clear that improved performance in extreme event quantification does not necessarily mean greater performance in classification, as measured by the area under the precision recall curve. However, this result requires distinct thresholds for what is considered an extreme event under f and \hat{f} . Thus, the area under the precision recall curve interprets \hat{f} not as an interpolation of f but as a distinct metric used for classification.

The “best” loss function for a particular task will depend on the problem, dataset, and goals. Of those studied in this work, L_{MSE} and L_{RE} have the advantage of not requiring an estimate of p_y . The extent to which this proves troublesome will depend on the dimension of y , sample size, and on the distribution p_y . Tail events are most accurately quantified by L_{OW} , though at the cost of poor performance on events in the core of p_y . This is partially mitigated by L_{AOW} . Networks used in any application setting should be selected via cross-validation using a problem-appropriate metric.

An obvious challenge in the use of loss functions which include p_y is that the number of samples required for density approximation grows geometrically in the dimension of y . In this work we have focused on the case where y is scalar valued and thus avoided the issue. Problems in higher dimensions may require approximating density through some observable or via distance to neighbors along a low dimensional manifold via technique like diffusion maps [12], if such a manifold exists. The relative entropy loss does not require approximation of p_y , but does equate importance of a particular sample with its exponentiated magnitude. Thus, it may not be an ideal choice when the rare events of interest are not substantially different in magnitude from the core of the distribution. Further investigation of such problems would be an interesting research direction, but we consider it to be outside the scope of this work.

The proposed extreme event loss functions may be a poor choice for certain classes of prediction problems, even if the data has extreme events. Consider for example the task of learning a dynamic model for a system with extreme events using short time series or velocity data as in [13, 53]. In this scenario, errors in the prediction phase are compounded and the higher error rate in the core of the distribution will render long term forecasts less accurate. This is in contrast to models trained with L_{MSE} where error accumulation will likely be more focused on rare events.

There exists a wide spectrum of functional forms used in machine learning that allows researchers to select or construct models they deem fit for a particular task. Less attention is paid to the choice of functional indicating the performance of those models, though some works have considered this question and proved its importance [43, 13]. This may be due to the wide effectiveness of the often used mean square error, and its clear motivation as a maximum likelihood estimate given Gaussian error. However, as we have shown, performance on certain tasks is significantly improved with tailored loss functions. The present manuscript extends this important research direction.

Funding:

This material is based upon work supported by the National Science Foundation (1902972), the Office of Navy Research (N00014-21-1-2357), and the Airforce Office of Scientific Research (MURI FA9550-21-1-0058). Simulation of the flow around a square cylinder was run using the Extreme Science and Engineering Discovery Environment (XSEDE) [52] allocation TG-MTH210003. XSEDE is supported by National Science Foundation under Grant No. ACI-1548562.

Appendix A: Details on datasets

Numerical simulations of the Kolmogorov flow and flow around a square cylinder were performed using Nek5000 [42], an open source Fortran based spectral element code for incompressible flow [41]. Time integration was carried out using a second order semi-implicit scheme described in [22]. Simulations were stabilised with a small degree of filtering as described in [21]. Mesh was generated using gmsh [23]. The Kolmogorov flow dataset uses 144 elements of order 7 and the flow around a square cylinder uses 1728 elements of order 7. The domain for the square cylinder flow extended from $x = -12$ to $x = 30$ and $|y| \leq 12$ with the square having sidelength 1 centered at the origin. Geometry and case files for rerunning simulations as well as files with numerical values of output data used in this work are available on GitHub at https://github.com/snagcliffs/EE_loss.

The mean and standard deviation of the energy dissipation for the Kolmogorov flow dataset used in this work are 0.116065 and 0.037559, which are respectively within 0.64% and 2.1% of those reported in [19]. The grid for the flow around a square cylinder was validated via comparison to a short simulation using a finer mesh. The mean drag coefficient on the interval $t \in [200, 2000]$ using 1728 spectral elements was found to be approximately 1.34% different from that using 4480 elements, which was deemed sufficiently resolved for the purpose of this work.

Appendix B: Neural networks details

Neural networks used in this manuscript were implemented in Python using the tensorflow library [1] as well as the Numba library [32]. Gaussian process models for p_y were initially trained using GPy [25], a Python library for Gaussian processes. Learned parameters were subsequently used to build non-trainable Gaussian process models in tensorflow, allowing for their use with neural network optimization tools. Source code for neural networks and scripts used to train examples used in this work are available on GitHub at https://github.com/snagcliffs/EE_loss.

Networks were built using a combination of dense layers and long-short-term-memory (LSTM) layers [29]. Networks for the Kolmogorov flow had 3 dense layers of size (4,8,16), followed by an LSTM layer with (32) units, followed by dense layers of size (16,8,4). Networks for the square cylinder had 3 dense layers of size (4,8,16), followed by an LSTM layer with (16) units, followed by dense layers of size (16,8,4). In each case the swish activation function was used [46].

Neural networks are widely known to be prone to over fitting. To mitigate this, we used early stopping [40] and noise injection in the training data. Data was separated into disjoint temporally contiguous sets for training, validation, and testing using a (50/10/40%) split and training was stopped when validation set loss failed to yield a new minimum for 5 consecutive epochs. Noise injection has been shown to improve generalization performance [2, 3], which may be due in part to its relation to Tikhonov regularization [4]. During the training procedure, we sampled random noise from a Gaussian distribution having standard deviation equal to 10% of the data and added it to both the neural network input.

References

- [1] Martín Abadi and et. al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from [tensorflow.org](https://www.tensorflow.org).
- [2] Guozhong An. The effects of adding noise during backpropagation training on a generalization performance. *Neural computation*, 8(3):643–674, 1996.

- [3] Anna Asch, Ethan Brady, Hugo Gallardo, John Hood, Bryan Chu, and Mohammad Farazmand. Model-assisted deep learning of rare extreme events from partial observations. *arXiv preprint arXiv:2111.04857*, 2021.
- [4] Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.
- [5] Antoine Blanchard and Themistoklis Sapsis. Bayesian optimization with output-weighted optimal sampling. *Journal of Computational Physics*, 425:109901, 2021.
- [6] Patrick J Blonigan, Mohammad Farazmand, and Themistoklis P Sapsis. Are extreme dissipation events predictable in turbulent fluid flows? *Physical Review Fluids*, 4(4):044606, 2019.
- [7] MP Brenner, JD Eldredge, and JB Freund. Perspective on machine learning for advancing fluid mechanics. *Physical Review Fluids*, 4(10):100501, 2019.
- [8] Noah D Brenowitz and Christopher S Bretherton. Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12):6289–6298, 2018.
- [9] Steven L Brunton, Bernd R Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020.
- [10] Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019.
- [11] Joseph T Chang and David Pollard. Conditioning as disintegration. *Statistica Neerlandica*, 51(3):287–317, 1997.
- [12] Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- [13] NAK Doan, W Polifke, and L Magri. Short-and long-term predictions of chaotic flows and extreme events: a physics-constrained reservoir computing approach. *Proceedings of the Royal Society A*, 477(2253):20210135, 2021.
- [14] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *Annual Review of Fluid Mechanics*, 51:357–377, 2019.
- [15] Kristian Dysthe, Harald E Krogstad, and Peter Müller. Oceanic rogue waves. *Annu. Rev. Fluid Mech.*, 40:287–310, 2008.
- [16] David R Easterling, J L Evans, P Ya Groisman, Thomas R Karl, Kenneth E Kunkel, and P Ambenje. Observed variability and trends in extreme climate events: a brief review. *Bulletin of the American Meteorological Society*, 81(3):417–426, 2000.
- [17] Bradley Efron. Nonparametric standard errors and confidence intervals. *Canadian Journal of Statistics*, 9(2):139–158, 1981.
- [18] Mohammad Farazmand. An adjoint-based approach for finding invariant solutions of navier–stokes equations. *Journal of Fluid Mechanics*, 795:278–312, 2016.
- [19] Mohammad Farazmand and Themistoklis P Sapsis. A variational approach to probing extreme events in turbulent dynamical systems. *Science advances*, 3(9):e1701533, 2017.
- [20] Mohammad Farazmand and Themistoklis P Sapsis. Extreme events: Mechanisms and prediction. *Applied Mechanics Reviews*, 71(5), 2019.
- [21] Paul Fischer and Julia Mullen. Filter-based stabilization of spectral element methods. *Comptes Rendus de l’Académie des Sciences-Series I-Mathematics*, 332(3):265–270, 2001.
- [22] PF Fischer. Implementation considerations for the oifs/characteristics approach to convection problems. *Argonne National Laboratory*, 2003.
- [23] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.
- [24] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [25] GPy. GPy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy>, since 2012.
- [26] Abhinav Gupta and Pierre FJ Lermusiaux. Neural closure models for dynamical systems. *Proceedings of the Royal Society A*, 477(2252):20201004, 2021.
- [27] Stephen Guth and Themistoklis P Sapsis. Machine learning predictors of extreme events occurring in complex dynamical systems. *Entropy*, 21(10):925, 2019.

- [28] Haibo He and Edwardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.
- [29] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [30] Christopher Irrgang, Niklas Boers, Maike Sonnewald, Elizabeth A Barnes, Christopher Kadow, Joanna Staneva, and Jan Saynisch-Wagner. Towards neural earth system modelling by integrating artificial intelligence in earth system science. *Nature Machine Intelligence*, 3(8):667–674, 2021.
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [32] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.
- [33] François M Longin. The asymptotic distribution of extreme stock market returns. *Journal of business*, pages 383–408, 1996.
- [34] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- [35] Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1):1–10, 2018.
- [36] AJ Majda, DW McLaughlin, and EG1431687 Tabak. A one-dimensional model for dispersive wave turbulence. *Journal of Nonlinear Science*, 7(1):9–44, 1997.
- [37] Andrew McHutchon. Differentiating gaussian processes. *Cambridge (ed.)*, 2013.
- [38] Michele Milano and Petros Koumoutsakos. Neural network modeling for near wall turbulent flow. *Journal of Computational Physics*, 182(1):1–26, 2002.
- [39] Mustafa A Mohamad and Themistoklis P Sapsis. Sequential sampling strategy for extreme event statistics in nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 115(44):11138–11143, 2018.
- [40] Nelson Morgan and Hervé Bourlard. Generalization and parameter estimation in feedforward nets: Some experiments. *Advances in neural information processing systems*, 2:630–637, 1989.
- [41] Anthony T Patera. A spectral element method for fluid dynamics: laminar flow in a channel expansion. *Journal of computational Physics*, 54(3):468–488, 1984.
- [42] James W. Lottes Paul F. Fischer and Stefan G. Kerkemeier. nek5000 Web page, 2008. <http://nek5000.mcs.anl.gov>.
- [43] Di Qi and Andrew J Majda. Using machine learning to predict extreme events in complex systems. *Proceedings of the National Academy of Sciences*, 117(1):52–59, 2020.
- [44] Elizabeth Qian, Boris Kramer, Benjamin Peherstorfer, and Karen Willcox. Lift & learn: Physics-informed machine learning for large-scale nonlinear dynamical systems. *Physica D: Nonlinear Phenomena*, 406:132401, 2020.
- [45] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [46] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [47] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.
- [48] Stephan Rasp, Michael S Pritchard, and Pierre Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39):9684–9689, 2018.
- [49] Themistoklis P Sapsis. Output-weighted optimal sampling for bayesian regression and rare event statistics using few samples. *Proceedings of the Royal Society A*, 476(2234):20190834, 2020.
- [50] Themistoklis P Sapsis. Statistics of extreme events in fluid flows and waves. *Annual Review of Fluid Mechanics*, 53:85–111, 2021.
- [51] David Siegmund. Importance sampling in the monte carlo study of sequential tests. *The Annals of Statistics*, pages 673–684, 1976.

- [52] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazelwood, Scott Lathrop, Dave Lifka, Gregory D Peterson, et al. Xsede: accelerating scientific discovery. *Computing in science & engineering*, 16(5):62–74, 2014.
- [53] Zhong Yi Wan, Pantelis Vlachas, Petros Koumoutsakos, and Themistoklis Sapsis. Data-assisted reduced-order modeling of extreme events in complex dynamical systems. *PloS one*, 13(5):e0197704, 2018.
- [54] Larry Wasserman. *All of nonparametric statistics*. Springer Science & Business Media, 2006.
- [55] PK Yeung, XM Zhai, and Katepalli R Sreenivasan. Extreme events in computational turbulence. *Proceedings of the National Academy of Sciences*, 112(41):12633–12638, 2015.