

# Discovering and forecasting extreme events via active learning in neural operators

Ethan Pickering (✉ [pickering@mit.edu](mailto:pickering@mit.edu))

Massachusetts Institute of Technology <https://orcid.org/0000-0002-4485-6359>

George Karniadakis

Brown University <https://orcid.org/0000-0002-9713-7120>

Themistoklis Sapsis

Massachusetts Institute of Technology

---

## Article

### Keywords:

**Posted Date:** April 27th, 2022

**DOI:** <https://doi.org/10.21203/rs.3.rs-1527132/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

# Discovering and forecasting extreme events via active learning in neural operators

Ethan Pickering<sup>1,\*</sup>, George Em Karniadakis<sup>2</sup>, and Themistoklis P. Sapsis<sup>1,\*</sup>

<sup>1</sup>Department of Mechanical Engineering, Massachusetts Institute of Technology,  
Cambridge, MA 02138, USA

<sup>2</sup>Division of Applied Mathematics, Brown University, Providence, RI 02906, USA

April 22, 2022

## Abstract

Extreme events in society and nature, such as pandemic spikes or rogue waves, can have catastrophic consequences. Characterizing extremes is difficult as they occur rarely, arise from seemingly benign conditions, and belong to complex and often unknown infinite-dimensional systems. Such challenges render attempts at characterizing them as moot. We address each of these difficulties by combining novel training schemes in Bayesian experimental design (BED) with an ensemble of deep neural operators (DNOs). This model-agnostic framework pairs a BED scheme that actively selects data for quantifying extreme events with an ensemble of DNOs that approximate infinite-dimensional nonlinear operators. We find that not only does this framework clearly beat Gaussian processes (GPs) but that 1) shallow ensembles of just *two* members perform *best*; 2) extremes are uncovered regardless of the state of initial data (i.e. with or without extremes); 3) our method eliminates “double-descent” phenomena; 4) the use of batches of suboptimal acquisition points compared to step-by-step global optima does not hinder BED performance; and 5) Monte Carlo acquisition outperforms standard minimizers in high-dimensions. Together these conclusions form the foundation of an AI-assisted experimental infrastructure that can efficiently infer and pinpoint critical situations across many domains, from physical to societal systems.

## 1 Introduction

The grand challenge of predicting disasters remains an extremely difficult and unsolved problem [1]. Disasters, such as pandemic spikes, wildfires or rogue waves<sup>1</sup>, are uniquely challenging to quantify. This is because they are both *rare* and arise from an *infinite* set of physical conditions [4]. The proposition of predicting extremes is analogous to *finding a catastrophic needle in an infinite-dimensional haystack*. This calls for methods that can both *discover* extreme events and encode *physical* phenomena into their modeling strategy. We present a Bayesian-inspired experimental design (BED) approach, described in detail in figure 5, that addresses both challenges by combining a probabilistic “discovery” algorithm [5, 6] with a deep neural operator designed to approximate physical systems [7].

Discovery of extremes is often simplified by distilling complex systems to their governing input variables and relevant output variables. Within this interpretation, quantification of extremes has historically taken the form of importance sampling, which uses a biasing distribution to identify regions of the input space that exhibit extreme values [8, 9]. Unfortunately, these techniques often require additional and challenging considerations for accurate results [10, 11, 12] and are static, lacking an ability to adjust to new information gained through experiments. BED provides a dynamic approach that learns from acquired data before

---

\*Co-corresponding authors: pickering@mit.edu, sapsis@mit.edu

<sup>1</sup>Rogue waves are rare, giant waves that pose a danger to ships and offshore structures. The largest wave on record was 25.6 meters and hit the Draupner oil platform in the North Sea on January 1, 1995 [2]. The “most extreme” wave, three times the size of surrounding waves, was 19.5 meters and observed off the coast of Ucluelet, British Columbia, on November 17, 2020 [3]

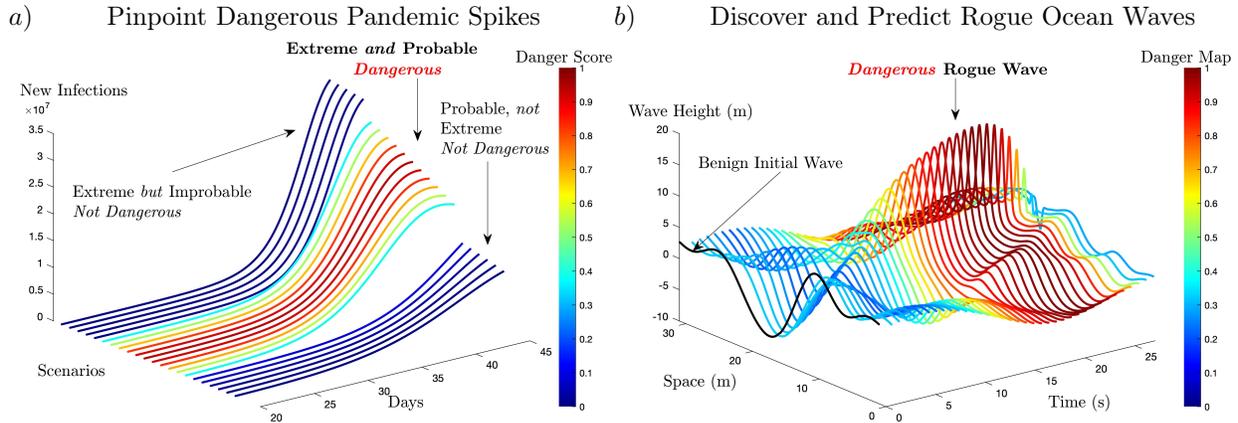


Figure 1: **Inference of diverse extreme phenomena, from pandemic spikes to rogue ocean waves.** Our framework pinpoints the most dangerous (i.e., probable and extreme) pandemic scenarios (left) and discovers rogue waves (right). The most dangerous pandemic scenarios are pinpointed by inferring the number of new infections in time for a plurality of infection rate hypotheses. See figure 2 for more details. Rogue waves are discovered and quantified for future prediction by uncovering the probable wave conditions that non-linearly interact in time to emit rogue waves over three times their original size. We show one example of this phenomena here and refer to figure 4 for more details on discovering these waves.

38 selecting new and intriguing input-output data. However, BED has not been a popular choice in rare-event  
 39 characterization due to a lack of 1) surrogate models that accurately generalize infinite-dimensional systems,  
 40 and 2) appropriately defined acquisition functions for selecting extreme data.

41 Deep neural operators (DNO), such as DeepONet [7], are built specifically for handling infinite-dimensional  
 42 systems and provide the ideal surrogate model for characterizing extremes. Unlike other ML approaches,  
 43 such as Gaussian Process Regression (GP), which map parameterizations of physical phenomenon, DNOs  
 44 directly map *physical*, infinite-dimensional functions to *physical*, infinite-dimensional functions. This leads to  
 45 drastic improvements in generalization to unseen data in high dimensions. Additionally, the neural network  
 46 backbone of DNOs mean they are intrinsically amenable to big data, unlike GPs which scale as the third  
 47 power of data size [13, 14]. However, DNOs utility for Bayesian experimental design is an open question  
 48 as DNOs do not explicitly provide a measure of uncertainty. We propose and show the efficacy of using an  
 49 ensemble of DNOs for uncertainty quantification and BED. Although much of the literature is skeptical of  
 50 the generality of ensembles to provide uncertainty estimates, recent viewpoints, [15] and notably [16], have  
 51 argued that DNN ensembles provide a very good approximation of the posterior. Our results support this  
 52 perspective.

53 Appropriately defined acquisition functions for uncovering extreme behavior are just as critical as the  
 54 chosen surrogate model. Recently, [5], in concert with several other works [17], [18], and [6], introduced a  
 55 class of probabilistic acquisition functions specifically designed for discovering extreme events. By combining  
 56 statistics of the input space along with statistics deduced from the surrogate model, the method can account  
 57 for both the importance of the output relative to the input. This approach significantly reduces the number  
 58 of input samples required to characterize extreme phenomena.

59 Equipped with these acquisition functions and an ensemble of DNOs, we will search high-dimensional  
 60 systems to discover extreme rogue waves and pinpoint dangerous pandemic scenarios, the two classes of  
 61 representative problems we consider herein.

## 62 2 Results

63 Figure 1 presents the key implications of our results, diagnosing the most dangerous future pandemic scenarios  
 64 (left) and discovering unassuming waves that lead to dangerous rogue waves (right). In each case, different

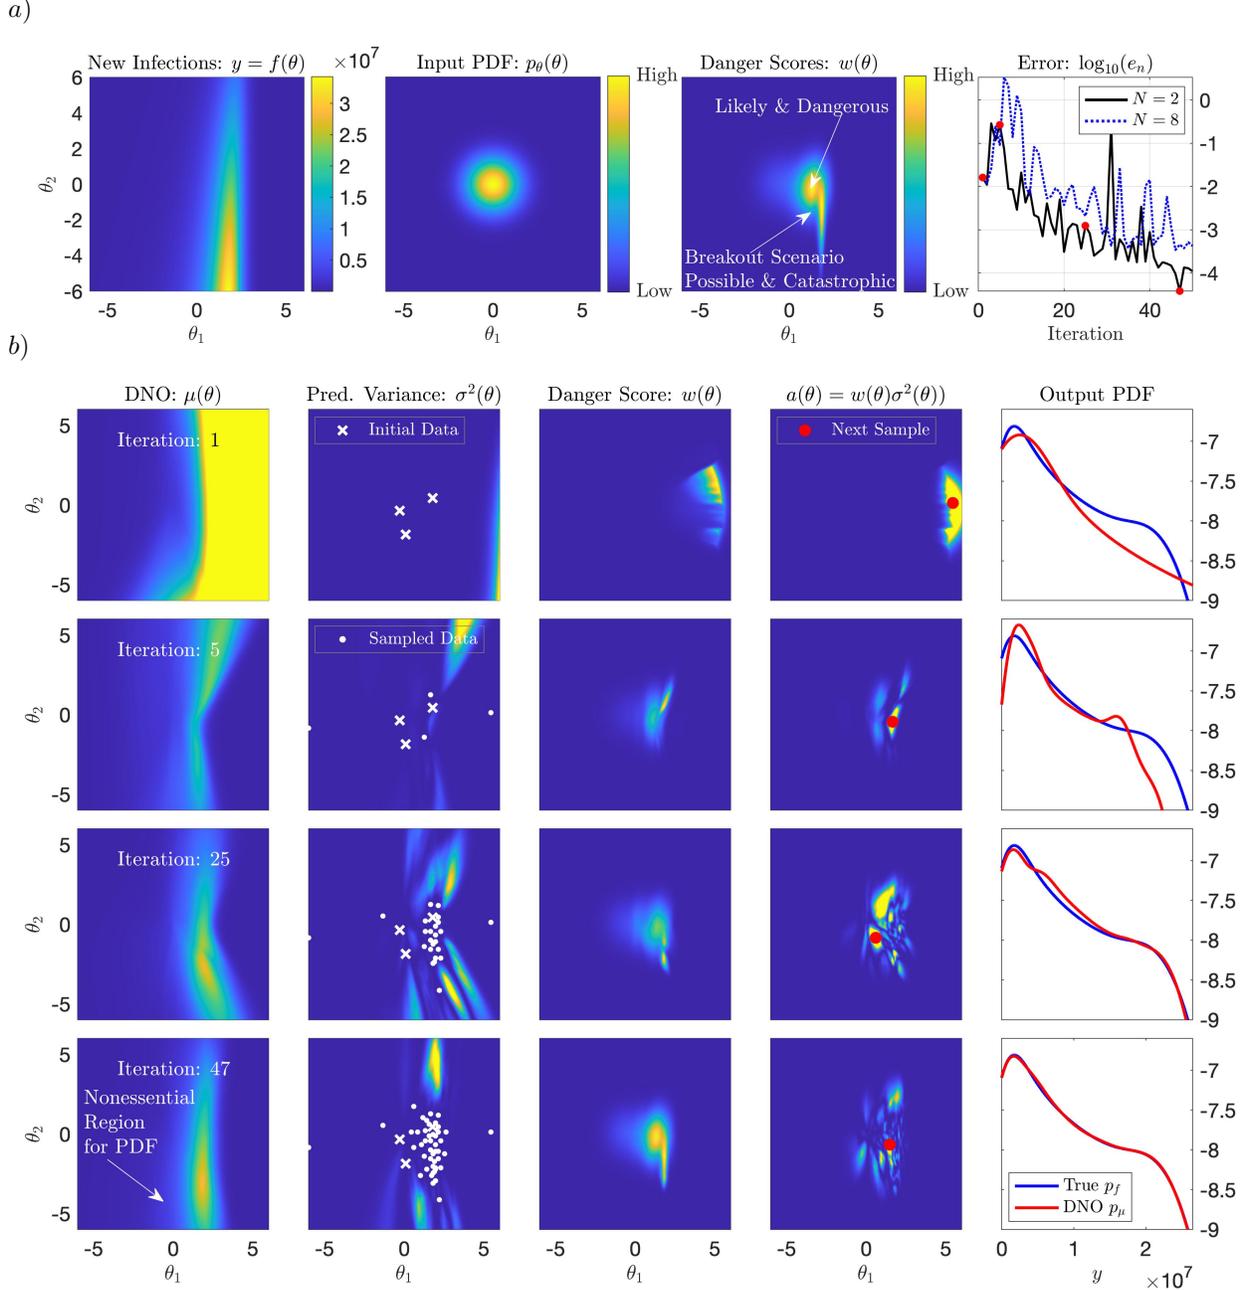


Figure 2: **Nailing the Tail: Accurate PDF and danger map convergence in 50 samples.** *a)* From left to right, the full deterministic response of new infections,  $f(\theta)$ , with respect to the two random parameters,  $\theta_1$  and  $\theta_2$ , the probability distribution of the random parameters  $\theta_1, \theta_2$ , the underlying danger scores,  $w(\theta)$ , and the  $\log_{10}$  of the log-pdf error for the experiment performed in *b)* using  $N = 2$ . The red circles indicate the iterations shown in *b)* as well as additional results for a case with  $N = 8$  ensembles. *b)* One experiment of the 2D stochastic SIR model using three initial points and iterated 50 times. The rows give the iteration number, 1, 5, 25, and 47 from top to bottom, respectively. The columns, from left to right, are the DNO approximation of the objective function,  $\mu(\theta)$ , given the training points (initial + samples), the points sampled (where initial points are white x and sampled points are white circles) in the 2D parameter space and the predictive variance  $\sigma^2(\theta)$ , the calculated danger scores  $w(\theta)$ , the acquisition values  $a(\theta)$  with the next acquisition point denoted by a red circle, and the DNO approximated and true output PDFs. Animation links presenting 100 iterations for  $N = 2$  ([here](#)) and  $N = 8$  ([here](#)).

65 scenarios are tied to a “danger score” or *likelihood ratio*, as proposed by [5],

$$w(\boldsymbol{\theta}) = \frac{p_{\boldsymbol{\theta}}(\boldsymbol{\theta})}{p_{\mu}(\mu(\boldsymbol{\theta}))}, \quad (1)$$

66 where  $p_{\boldsymbol{\theta}}$  is the probability density function (PDF) of the random variable input,  $\boldsymbol{\theta}$ , and  $p_{\mu}$  is the posterior  
67 PDF of the mean,  $\mu$ , conditioned on the input,

$$p_{\mu}(\mu) = \int_{\boldsymbol{\theta}} p_{\mu}(\mu|\boldsymbol{\theta})p_{\boldsymbol{\theta}}(\boldsymbol{\theta})d\boldsymbol{\theta}. \quad (2)$$

68 The likelihood ratio appropriately balances events that are likely and those that are extreme, hence it  
69 provides a danger score for any given event. As denoted for the pandemic model in figure 1, small danger  
70 scores are attributed to events that are either implausible or not extreme, while large danger scores relate  
71 to those that are both likely and extreme. However, any system’s danger score requires knowledge of the  
72 posterior mean,  $p_{\mu}$ , which is generally unknown and must be learned. Our approach efficiently learns this  
73 underlying distribution through dynamic application of the danger score with Bayesian experimental design  
74 and deep neural operators.

## 75 2.1 Pinpointing Dangerous Pandemic Scenarios

76 Figure 2 a) and b) demonstrate how the proposed framework efficiently learns the underlying output dis-  
77 tribution of infections for a stochastic pandemic model via dynamic and sequential application of danger  
78 scores. The pandemic model is a simple Susceptible, Infected, Recovered (SIR) model with a two-dimensional  
79 stochastic infection rate (see Appendix A). We assess success as the log-PDF error in figure 2 a) last column,  
80 between the true output distribution and the approximated distribution from the trained DNO, figure 2 b)  
81 last column. See Section 4.2.3 for details on computing the log-PDF error metric.

82 Despite an initialization of only three data points in the parameter space, our framework quickly identifies  
83 the key regions of dynamical relevance and accurately recovers the significant properties of the underlying  
84 system. Using an ensemble of just two DNOs (see Appendix D for DNO implementation and section 4.1.1 for  
85 our application of ensemble methods), the algorithm iteratively provides an estimation of the underlying map  
86 (for computing  $p_{\mu}$ ), variance,  $\sigma^2(\boldsymbol{\theta})$ , and a danger score,  $w(\boldsymbol{\theta})$ . Together, the danger score and variance create  
87 the acquisition function,  $a(\boldsymbol{\theta}) = w(\boldsymbol{\theta})\sigma^2(\boldsymbol{\theta})$ , that identifies the point within the parameter space with the  
88 greatest potential for learning the true output PDF. With the addition of each point, all fields dynamically  
89 change and bring the true and approximated output PDFs within greater agreement. By iteration 50, the  
90 danger scores have converged and the approximated output PDF is in error of less than  $10^{-3}$ . It is from this  
91 final danger score map that we derive the pandemic scenarios of figure 1 and note that this map includes both  
92 dangerous pandemic spike scenarios and a breakout scenario with catastrophic consequences. Additionally,  
93 figure 2 a) last column, shows that increasing the ensemble size to  $N = 8$  provides little to no advantage.

94 The critical aspect of this approach is the algorithm’s reduction of a large parameter space to local  
95 regions of danger. Only regions that provide significant contributions to the output PDF are considered.  
96 Iteration 47 in figure 2 b) underscores this behavior. The algorithm has accurately reconstructed the output  
97 PDF, yet it ignores a large region of high magnitude. This region, and the remaining unexplored regions,  
98 provides negligible information and is neglected. This property is crucial for all systems where resources for  
99 experiments or simulation are limited or costly and it permits a significant reduction in training/acquired  
100 data as system complexity and dimensionality increases.

## 101 2.2 Discovering and Predicting Rogue Waves

102 We now scale the method to higher dimensions, larger datasets, and on a significantly more complex and  
103 nonlinear system: a dissipative version of the nonlinear Schrödinger equation used to model rogue waves  
104 (Appendix B). Figure 3 a)-c) shows that the DNO-BED framework, using the likelihood-weighted uncertainty  
105 sampling (US-LW,  $a(\boldsymbol{\theta}) = w(\boldsymbol{\theta})\sigma^2(\boldsymbol{\theta})$ ), is superior for approximating the output PDF of the danger map (see  
106 section 4.2.3 for NLS danger map definition) than using GPs (detailed in Appendix C) or other common BED  
107 sampling strategies, such as uncertainty sampling (US,  $a(\boldsymbol{\theta}) = \sigma^2(\boldsymbol{\theta})$ ) and latin hypercube sampling (LHS).

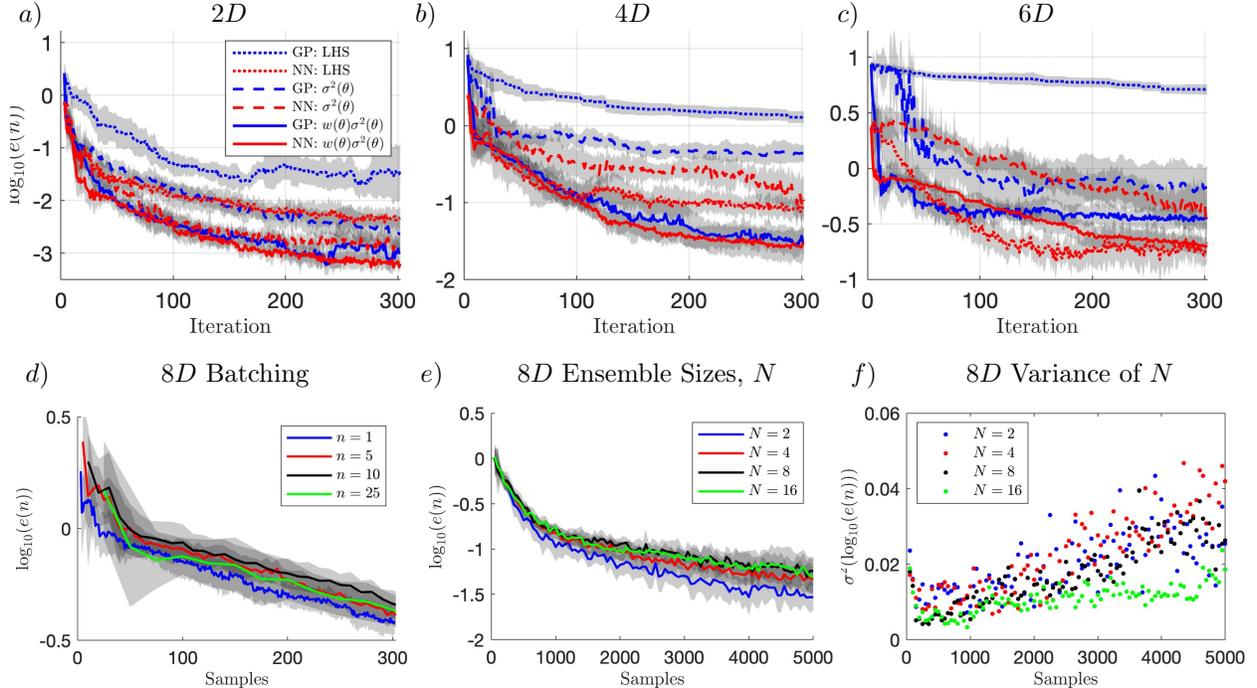


Figure 3: **Accelerated convergence with DNOs plus extreme acquisition functions regardless of dimensionality.** *a) – c)* Median log-PDF of 10 independent and randomly initialized experiments, where shaded regions denote one standard deviation for each case (i.e.  $\pm\sigma(\epsilon)$ ), considering three acquisition functions and GP and DNO surrogate models. *a) – c)* are for 2D, 4D and 6D initial conditions, while *d) – f)* represent results from 8D initial conditions. **Parallel acquisition and shallow ensembles bring computational efficiency without performance loss.** *d)* provides errors related to different batch sizes per iteration, *e)* gives performance differences relative to the DNO ensemble size (100 iterations at batch size  $n = 50$ ), and *f)* presents the variance in log-PDF error of *e)*.

108 This is highlighted as dimensionality (i.e., complexity, details on dimensionality provided in Appendix B) is  
 109 increased from 2D to 6D, where GPs begin to break down.

110 Figure 3 *d)* and *e)* provide two major implications for reducing the framework’s computational cost in  
 111 high dimensions: neither batching samples via parallel selection of multiple acquisition points, *d)*, nor use  
 112 of shallow ensembles of only  $N = 2$  members, *e)*, hinders performance. Computed at 8D, figure 3 *d)* shows  
 113 that regardless of choosing 1,5,10, or even 25 points per iteration, over 300 samples, does not result in a loss  
 114 of performance with our framework. Further, we show in Appendix F that the use of Monte Carlo methods  
 115 are substantially more efficient for identifying optimal acquisition points than standard PYTHON optimizers.  
 116 See Section 4.4 for details on batching implementation. These observations are a critical result for scaling  
 117 our framework, as more complex systems inevitably will require more data.

118 Unexpectedly, figure 3 *e)* shows that not only do *just two ensemble members,  $N = 2$ , perform well,*  
 119 *they consistently outperform larger ensembles,  $N > 2$ .* This result appears to disagree with the hypothesis  
 120 that a larger set of ensembles would provide uncertainty estimates with greater fidelity, leading to better  
 121 performance of our sequential search methods. Clearly, this is not the case from our results, yet neither  
 122 can it be that  $N = 2$  ensembles provide a predictive variance of greater fidelity than  $N = 16$ . Figure 3  
 123 *f)* permits both concepts to be true. It shows that the greater the ensemble size, the smaller the variance  
 124 between the error trajectories of independent experiments. This observation agrees with the idea that larger  
 125 ensembles lead to a higher-fidelity predictive variance, but that greater fidelity leads to *consistency* rather  
 126 than *performance* for this sequential search technique. We believe using small  $N$  imposes a greedy search, in  
 127 a similar fashion to Thompson sampling [19]. Regardless, the consistent observation that  $N = 2$  is not only  
 128 viable, but perhaps preferable, has massive implications for minimizing computational costs for ensemble  
 129 approaches.

### 130 2.2.1 Rogue Wave Discovery in 20D

131 Equipped with the computational advantages of  $N = 2$  ensemble members and large batch sizes ( $n_b = 50$ ),  
132 figure 4 *a) – b)* shows that even at 20D our approach can recover the true output PDF. The other acquisition  
133 functions not only perform poorly, but perform worse as more data is acquired. This phenomenon is known as  
134 “double descent”, and many researchers have observed this behavior throughout machine learning procedures,  
135 from classification to regression problems [20]. To our knowledge, this is the first observation of double  
136 descent arising from the tails of the log-PDF error. Double descent is a product of over-fitting. More data  
137 and greater complexity results in an over-parameterization of the provided data. Temporarily, this leads to  
138 inferior generalization until providing the surrogate model with sufficiently larger datasets.

139 *The proposed acquisition function clearly avoids double descent.* We believe this is because of the efficient  
140 selection of points that significantly contribute to the observed dynamics of the system. Unlike US-LW, the  
141 points chosen by LHS and US methods are not inherently important to recovering the true PDF and therefore  
142 induce misleading complexity to the underlying regression task. This observation further underscores the  
143 value of our acquisition function, as it systematically prevents over-fitting and unwarranted model complexity.

144 Figure 4 *a) – b)* also show that regardless of the origin of initial conditions, containing extremes (chosen  
145 by LHS) or without extremes (from the prior,  $p_\theta$ ), the method achieves similar error metrics. Figure 4 *c)*  
146 and *d)* present the physical manifestations of the 20D initial conditions that include extremes and do not.  
147 Despite these clearly different sets of initial data, the algorithm uncovers dangerous benign conditions that  
148 lead to extremes.

## 149 3 Discussion

150 The discovery, characterization, and ultimate prediction of extreme events remains a grand challenge across  
151 application domains. Presented in figure 5, we propose an AI-assisted, model-agnostic infrastructure that  
152 paves the way for efficient and robust inference of these critically important events. By combining deep neural  
153 operators, Bayesian experimental design, and tailored acquisition functions, we create a framework that can  
154 accurately discover, characterize, and predict extreme events in high-dimensional nonlinear systems ranging  
155 from societal to physical. We demonstrate the efficacy of this framework here by efficiently discovering and  
156 predicting rogue waves, reducing the number of samples from  $10^7$  to  $10^2$ , and accurately pinpointing the  
157 most dangerous pandemic spikes.

158 This approach is not only superior to other techniques, namely Gaussian Processes (GPs), but comes with  
159 several expected and unexpected, yet preferable, consequences. As expected, we show that the use of DNOs  
160 brings both greater generalization accuracy and a propensity for extremely large datasets. Unexpectedly,  
161 we find that shallow ensembles of only *two* ensemble members perform remarkably well and that the use  
162 of batches of optimal sample points acquired in parallel, compared to step-by-step, or serial, acquisition,  
163 does not hinder BED performance. We also find that Monte Carlo approaches for selecting these optimal  
164 points perform better than standard optimizers, which routinely fall victim to the highly non-convex nature  
165 of high-dimensional optimization. These results bring crucial real-world benefits to our approach. Batching  
166 permits parallel acquisition of data, while small ensembles keep computational costs tractable.

167 Our method also eliminates “double descent” phenomena and achieves accurate extreme event quantifi-  
168 cation *regardless* of the state of initial data. Both observations underscore the importance of sampling the  
169 *most informative data rather than more data*. Double descent is avoided as the acquisition function only  
170 selects points that significantly contribute to the observed dynamics of the system, while ignoring data that  
171 would bring added and unnecessary complexity. On the other hand, the method is robust to the initial data  
172 provided and quickly directs attention to the most informative samples whether or not extremes have been  
173 observed.

174 The framework also provides modularity for the chosen surrogate models and acquisition functions, so long  
175 as the parameter and function spaces are appropriately defined. Critically, we separate the two. Regression  
176 is performed in the functional space, while the search algorithm is performed in the parameter space via  
177 a forward DNO coupling (i.e., through a linear transfer function from  $\theta \rightarrow u(x)$ ). Any arbitrary neural  
178 architecture leveraging this distinction may be implemented, such as standard feed-forward NNs, Fourier  
179 Neural Operators [21], Convolutional NNs, Recurrent NNs, Long-Short Term Memory, among others. In  
180 fact, this work did not explicitly investigate the full value of DeepONet for operator learning for BED. Instead,

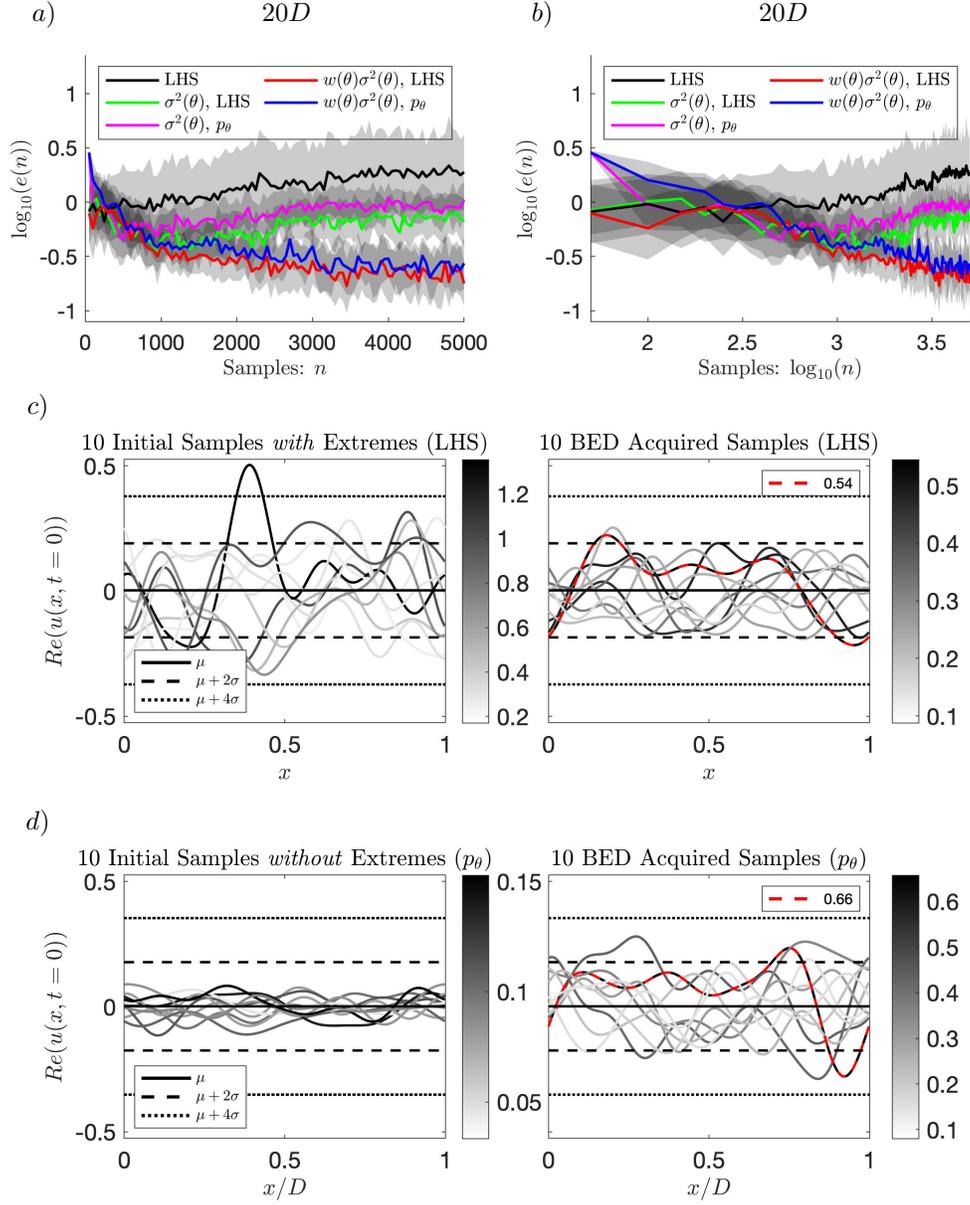


Figure 4: **Robustness to dimensionality and initial data, i.e., with or without extremes.** *a)* The log-PDF errors for three acquisition functions (LHS, US, and US-LW) for the 20D problem, with initial points sampled via LHS and the prior  $p_\theta$  for US and US-LW. *b)* provides the same results in log-log form. *c)* gives the real value of 10 complex LHS initial input functions  $u(x, t = 0)$ , shaded with their corresponding output value  $\|Re(u(x, t = \tau))\|_\infty$  (left) and 10 DNO+BED acquired functions, via US-LW at iteration 97 (right). *d)* gives 10 initial input functions sampled from the prior distribution  $p_\theta$  (left) and the 10 acquired function at iteration 74 (right).

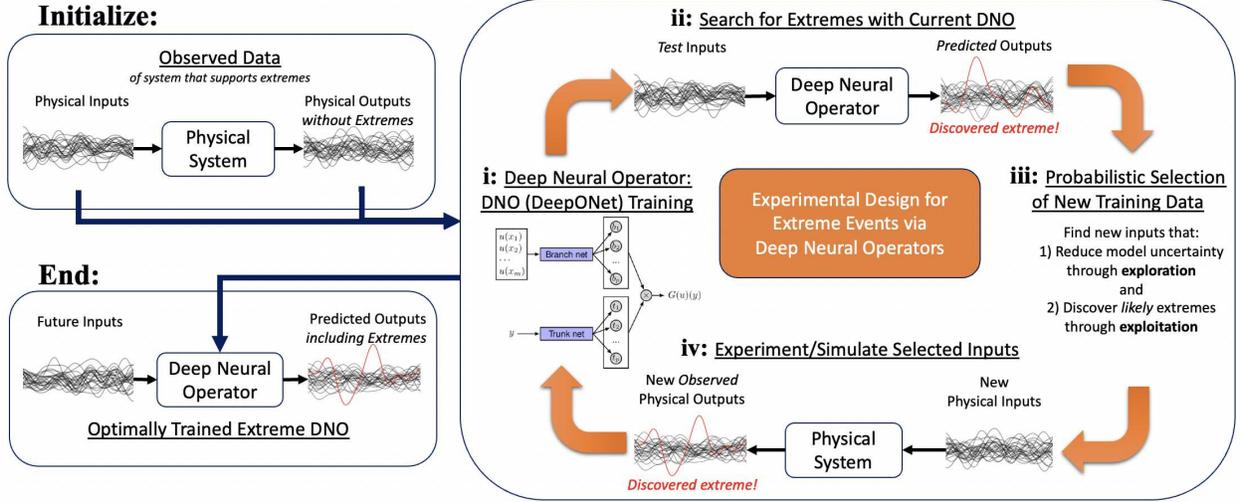


Figure 5: **Efficient and robust DNO+BED framework for discovering and quantifying extremes.** An overview of the proposed Bayesian experimental design framework with deep neural operators and novel acquisition functions for discovering extremes. *Initialize*) with a set of observed physical input-output pairs, retained in their functional form. *i)* Pass functions to an ensemble of DNOs to learn sparse representations of the underlying system. *ii)* Perform a fast Monte Carlo search of the DNO functional space for extremes. *iii)* Compute statistics over a Monte Carlo ensemble and select new input functions that both explore and exploit the space for extremes. *iv)* Evaluate proposed inputs on the underlying experiment or simulation, record outputs, and pass to *i)*. Repeat *i)-iv)* until statistics are converged or resources are depleted. *End*) with an optimally trained DNO that supports prediction of extreme events.

181 only the standard feed-forward branch NN was used. Using the complete operator machinery only requires  
 182 adjusting the parameters related to the operator and trunk. Here, we kept these parameters constant.  
 183 Finally, while we focus on Bayesian experimental design, slight adjustments to the choice of acquisition  
 184 function allow for Bayesian Optimization tasks [22] or for approaching other metrics of interest (e.g. mean  
 185 squared error).

186 In conclusion, we believe we have demonstrated an equation-agnostic framework that: (1) efficiently dis-  
 187 covers extreme events, (2) is computationally tractable, and (3) presents a straightforward implementation  
 188 on any arbitrary input-output system. This creates a unique opportunity for experimentalists and compu-  
 189 tationalists alike to investigate their systems for extreme behavior, whether that behavior originates from  
 190 societal or physical systems or has beneficial or catastrophic consequences.

## 191 4 Methods

192 Our approach to Bayesian experimental design is detailed in figure 5 and consists of two critical components,  
 193 *data selection criteria* and the *surrogate model*. Algorithm 1 formalizes the iterative steps taken in figure 5  
 194 for efficiently training a surrogate model with minimal data selection.

### 195 4.1 Surrogate Models

196 We test two surrogate models in the BED framework, GPs as the benchmark case, and DNOs (specifically  
 197 DeepONet) for greater scaling and generalization performance. While both GPs and DNOs provide the  
 198 same role in the framework, their implementation is fundamentally different. Whereas GPs are used to map  
 199 random parameter inputs  $\theta$  to an output  $f(\theta)$ , we use the DNOs to map an input function  $u = \theta\Phi$  to an  
 200 output function  $G(u = \theta\Phi)$ . In both cases,  $\theta$  is the only independent variable and is maintained to provide  
 201 a means of searching for extremes in the parameter space. However, the proposed regression task for DNOs

---

**Algorithm 1** Sequential search for active learning/training of DNOs and GPs .
 

---

- 1: **Input:** Number of iterations:  $n_{iter}$
  - 2: **Initialize:** Train GP/DNO on initial dataset of input-output pairs  
 GP:  $\mathcal{D}_0 = \{\theta_i, y_i = f(\theta_i)\}_{i=1}^{n_{init}}$       DNO:  $\mathcal{D}_0 = \{\theta_i, y_i = G(\theta_{u,i}\Phi(x_1, \dots, x_m), \theta_{z,i})\}_{i=1}^{n_{init}}$
  - 3: **for**  $n = 1$  **to**  $n_{iter}$  **do**
  - 4:    Select next point  $\theta_n$  by minimizing\* acquisition function  $a(\theta)$ :  
     \*Minimization using Monte Carlo, See Section 4.3 and Appendix F
  - $$\theta_n = \arg \min_{\theta \in \mathcal{X}} a(\theta; y), \mathcal{D}_{n-1}$$
  - 5:    Evaluate objective function at  $\theta_n$  and record  $y_n$
  - 6:    Augment dataset:  $\mathcal{D}_n = \mathcal{D}_{n-1} \cup \theta_n, y_n$
  - 7:    Retrain GP/DNO.
  - 8: **end for**
  - 9: **return** Final GP/DNO Model
- 

ensures that they perform their mapping in functional, or physical, space, rather than the parameter space. This distinction is foundational for our approach and success with DNOs for BED.

We detail both GPs and DeepONet in detail in Appendices C and D, respectively, and only discuss our approach for quantifying the predictive variance,  $\sigma^2(\theta)$ , for DNOs using ensembles.

#### 4.1.1 Ensemble of Neural Networks for Uncertainty Quantification

Although neural network architectures are attractive for approximating nonlinear regression tasks, their complexity rids them of analytical expressions. This does not allow for a traditional Bayesian treatment of uncertainty in the underlying surrogate model – a key property present for GP regression (see equation 22). Knowledge of the uncertainty of a surrogate model allows one to target model deficiencies as seen in the parameter space. This means that choosing an appropriate method for quantifying the uncertainty is a crucial and key component to active learning or BED. There are several techniques for quantifying uncertainty in neural networks and we provide a brief description of these in Appendix D.1 and focus only on ensemble methods.

Ensemble approaches have been used extensively throughout the literature [23, 24] and despite their improved results for identifying the underlying tasks at hand [25], their utility for quantifying uncertainty in a model remains a topic of debate. There are several approaches for creating ensembles. These include random weight initialization[24], different network architectures (including activation functions), data shuffling, data augmentation, bagging, bootstrapping, and snapshot ensembles [26, 27, 28] among others. Here we employ random weight initialization, a technique found to perform similarly or better than BNN approaches (Monte Carlo Dropout and Probabilistic Backpropagation) for evaluation accuracy and out-of-distribution detection for both classification and regression tasks [24]. As stated earlier, much of the field is skeptical of the generality of ensembles to provide rigorous uncertainty estimates. However, recent studies, such as [15] and specifically [16], have argued that DNN and DNO ensembles provide reasonable, if not superior, approximations of the posterior. Finally, the straightforward implementation of the randomly initialized weights motivates our choice, as it makes the adoption of these techniques far more probable.

We train  $N$  randomly weight-initialized DNO models, each denoted as  $\tilde{G}_n$ , that find the associated solution field  $y$  for inputs  $u$  and  $z$ . This allows us to then determine the point-wise variance of the models as

$$\sigma^2(u, z) = \frac{1}{(N - 1)} \sum_{n=1}^N (\tilde{G}_n(u)(z) - \overline{\tilde{G}(u)(z)})^2, \tag{3}$$

where  $\overline{\tilde{G}(u)(z)}$  is the mean solution of the model ensemble. Finally, we must adjust the above representation to match the description for BED. In the case of traditional BED and GPs, the input parameters,  $\theta$ , represent the union of two set of parameters,  $\theta_u$  and  $\theta_z$ . The parameters  $\theta_u$  typically represent coefficients to a set of functions that represent a decomposition of a random function  $u = \theta_u \Phi(x_1, \dots, x_m)$ , while  $\theta_z = z$  represent

non-functional parameters of the operator. Thus, the DNO description for UQ may be recast as:

$$\sigma^2(\boldsymbol{\theta}) = \sigma(\boldsymbol{\theta}_u \cup \boldsymbol{\theta}_z) = \frac{1}{(N-1)} \sum_{n=1}^N (\tilde{G}_n(\boldsymbol{\theta}_u \Phi(x_1, \dots, x_m))(\boldsymbol{\theta}_z) - \overline{\tilde{G}(\boldsymbol{\theta}_u \Phi(x_1, \dots, x_m))(\boldsymbol{\theta}_z)})^2. \quad (4)$$

## 4.2 Data Selection: Acquisition Functions

The acquisition function is the key component of the sequential search algorithm, as it guides algorithm 1 in exploring the input/parameter space and determines points at which the objective function is to be queried. Because of the lack of a closed analytical form of DNOs, we only consider two acquisition functions used previously with GPs on several test cases in [5]. The two functions we are interested in are the commonly used uncertainty sampling and the output-weighted uncertainty sampling proposed by [5]. The novelty here is that we apply them explicitly to DNOs (DeepONet) and present the advantages of DNOs compared to GPs as we apply them to a complex and high-dimensional problem.

### 4.2.1 Uncertainty Sampling

Uncertainty sampling (US) is one of the most broadly used active sampling techniques and identifies the point where the predictive variance is the greatest,

$$a_{US}(\boldsymbol{\theta}) = \sigma^2(\boldsymbol{\theta}). \quad (5)$$

Uncertainty sampling, also known as the active-learning-MacKay (ALC) algorithm [29], imposes an sequential search that evenly distributes uncertainty over the input space as it gains data. The popularity of US is due to three qualities: ease of implementation, inexpensive evaluation (for small datasets with GPs), and analytic gradients, the last of which permits the use of gradient-based optimizers.

### 4.2.2 Likelihood-Weighted Acquisition Functions

There are several “extreme event” likelihood-weighted acquisition functions that could be explored, as proposed by [6], but we elect to only test the likelihood-weighted uncertainty sampling (US-LW) acquisition function because of its simplicity in implementation. For US-LW, we augment the US sampling acquisition function with the previously described danger scores to give

$$a_{US-LW}(\boldsymbol{\theta}) = w(\boldsymbol{\theta})\sigma^2(\boldsymbol{\theta}), \quad (6)$$

such that both *highly uncertain* and *highly dangerous* regions are sampled.

To compute  $w(\boldsymbol{\theta})$ , we note that the approximated output PDF,  $p_\mu(\mu)$  is approximated via a kernel density estimator with  $n = 10^6$  test points. For the DNO cases, we chose to compute this with only the first ensemble member,  $\mu = \tilde{G}_1$ , to reduce computational costs. Similar to the ensemble results for  $N = 2$ , using only one ensemble member is akin to using Thompson sampling and performs without reduction in performance.

### 4.2.3 Danger Maps and Log-PDF Error Metrics

To test the ability of the DNO and GP Bayesian-inspired sequential algorithms to quantify extremes, we define a “Danger Map” for the pandemic scenarios and rogue waves. The rogue wave danger map is defined as,

$$f(\boldsymbol{\theta}) = \|Re(u(x, t = T; \boldsymbol{\theta}))\|_\infty, \quad (7)$$

where  $T = 20$ , while the pandemic danger map is,

$$f(\boldsymbol{\theta}) = I(t = T; \boldsymbol{\theta}), \quad (8)$$

where  $T = 45$  days. For each case, we then select a random set, chosen through Latin hypercube sampling, of  $10^5$  test points,  $\Theta \in \mathbb{R}^{d \times 10^5}$  and use a kernel density estimator to compute the true PDF  $p_f(y)$ , computed as

$$p_f(y) = \int_{\Theta} p_f(y|\Theta)p_{\Theta}(\Theta)d\Theta, \quad (9)$$

262 and the model’s prediction  $p_{\mu_n}(y)$  for the same test points,  $\Theta$ . To determine whether the testing data  
 263 appropriately identifies extremes, we compute the log-PDF error

$$e(n) = \int |\log_{10} p_{\mu_n}(y) - \log_{10} p_f(y)| dy. \quad (10)$$

### 264 4.3 Monte Carlo Minimization of Acquisition Functions

265 In our experiments, we consistently observe that acquisition points found through minimizers using gradient  
 266 descent are not globally optimal. Instead, Monte Carlo evaluation of the DNOs and GPs consistently find  
 267 superior optima. This is chiefly because of the non-convexity of the acquisition function. We may recall the  
 268 highly non-convex behavior of the 2D acquisition fields in figure 2 b), even with as little as  $\approx 5$  samples. This  
 269 non-convex nature emits many local minima that requires many initial search points to provide confidence  
 270 that the chosen optima is nearly global. As the minimizer progresses for each iteration, it must call the  
 271 DNO or GP, whereas a Monte Carlo approach may efficiently evaluate all points in one vector operation.  
 272 This means that for the same computation time as the minimizers, Monte Carlo sampling may evaluate a  
 273 substantially larger distribution of query points and return acquisition points with superior scores than that  
 274 of the optimizer. In Appendix F we show that acquisition scores found via Monte Carlo at  $20D$  consistently  
 275 outperform minimizers for similar computation times.

### 276 4.4 Experiment Batching

277 As systems become more complex, additional experiments/data are required to reduce errors for higher  
 278 dimensional cases, as observed in figure 3. Considering many experiments can be conducted in parallel,  
 279 we ask whether choosing multiple local minima of the acquisition function presents marginally reduced  
 280 performance than a purely sequential search. This is especially critical for situations where experimental  
 281 time is more costly than additional setups (e.g. protein or genetic design).

282 The purpose of batching is to find *multiple regions of local optima* of the acquisition function, rather than  
 283 finding *several optima in the same region*. To impose this idea, we create a constraint that no acquisition  
 284 point may reside closer than a distance  $r_{\min}$  to each other. We define  $r_{\min}$  as a fraction of the maximum  
 285 euclidean distance of the space being sampled,

$$r_{\min} = r_l \left( \sum_{d=1}^D (\theta_{d,+} - \theta_{d,-})^2 \right)^{1/2}, \quad (11)$$

286 where  $\theta_{d,+}$  and  $\theta_{d,-}$  are the maximum and minimum domain bounds of each parameter dimension  $d$  and  $r_l$   
 287 is the user defined percentage. In this work, we chose a static  $r_l = 0.025$ , but dynamic values based on the  
 288 packing of the parameter space would be an intriguing direction for increasing the efficacy of this approach.  
 289 Imposing this constraint requires an iterative processing of the acquisition scores, detailed in Algorithm 2.  
 For the batching applied in this study, each case uses a Monte Carlo querying of  $n_q = 10^6$  points.

---

#### Algorithm 2 Sequential selection of points for batch sampling

---

- 1: **Evaluate** acquisition function for  $n_q$  query points.
  - 2: **for** Acquisition points smaller than batch size  $n_a < n_b$
  - 3:   **Choose** the minimum score,  $\min(a) = a(\theta_c)$ , from  $n_q$  points.
  - 4:   **Augment**  $\Theta_a$  with chosen point,  $\theta_c$ .
  - 5:   **Compute** the distances,  $r$ , between the chosen point and the remaining query points.
  - 6:   **Eliminate** all points from  $n_q$  where  $r < r_{\min}$ .
  - 7: **end for**
  - 8: **return** Points for the next experiment:  $\Theta_a$
- 

290

### 291 Data Availability

292 Data is computed using the code shared below and all data plotted throughout this manuscript may be  
 293 found [here](#). If accepted, all data will be shared via an open source link.

## 294 Code Availability

295 Code pertaining to the sequential discovery algorithm of both the SIR and NLS problem is provided via  
296 this shareable [link](#). If accepted, code will be shared via an open source GitHub repository. DeepONet code  
297 framework can be found within the `deepxde` package on [GitHub](#).

## 298 Acknowledgments

299 The authors acknowledge support from DARPA under the programs of Dr. Jiangying Zhou (Grant No.  
300 HR00112110002) and Dr. Joshua Elliott (Grant No. HR00112290029) awarded to MIT.

## 301 References

- 302 [1] Creating a disaster resilient america: Grand challenges in science and technology. *National Academies*  
303 *Press*, 2005.
- 304 [2] O. E. Hansteen, H. P. Jostad, and T. I. Tjelta. Observed platform response to a monster wave. In *Field*  
305 *Measurements in Geomechanics*, pages 73–86, 2003.
- 306 [3] J. Gemmrich and L. Cicon. Generation mechanism and prediction of an observed extreme rogue wave.  
307 *Scientific Reports*, 12(1):1–10, 2022.
- 308 [4] T. P. Sapsis. Statistics of extreme events in fluid flows and waves. *Annual Review of Fluid Mechanics*,  
309 53:85–111, 2021.
- 310 [5] Antoine Blanchard and Themistoklis Sapsis. Output-weighted optimal sampling for bayesian exper-  
311 imental design and uncertainty quantification. *SIAM/ASA Journal on Uncertainty Quantification*,  
312 9(2):564–592, 2021.
- 313 [6] A. Blanchard and T. P. Sapsis. Bayesian optimization with output-weighted optimal sampling. *Journal*  
314 *of Computational Physics*, 425:109901, 2021.
- 315 [7] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via DeepONet  
316 based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229,  
317 2021.
- 318 [8] H. Kahn and A. W. Marshall. Methods of reducing sample size in Monte Carlo computations. *Journal*  
319 *of the Operations Research Society of America*, 1(5):263–278, 1953.
- 320 [9] M. Shinozuka. Basic analysis of structural safety. *Journal of Structural Engineering*, 109(3):721–740,  
321 1983.
- 322 [10] G. Dematteis, T. Grafke, and E. Vanden-Eijnden. Extreme event quantification in dynamical systems  
323 with random components. *SIAM/ASA Journal on Uncertainty Quantification*, 7(3):1029–1059, 2019.
- 324 [11] F. Uribe, I. Papaioannou, Y. M. Marzouk, and D. Straub. Cross-entropy-based importance sampling  
325 with failure-informed dimension reduction for rare event simulation. *SIAM/ASA Journal on Uncertainty*  
326 *Quantification*, 9(2):818–847, 2021.
- 327 [12] S. Wahal and G. Biros. Bimc: The Bayesian Inverse Monte Carlo method for goal-oriented uncertainty  
328 quantification. Part I. *arXiv preprint arXiv:1911.00619*, 2019.
- 329 [13] E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. *Advances in Neural*  
330 *Information Processing Systems*, 18:1257, 2006.
- 331 [14] M. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelli-*  
332 *gence and Statistics*, pages 567–574. PMLR, 2009.

- 333 [15] E. Pickering and T. P. Sapsis. Structure and distribution metric for quantifying the quality of uncer-  
334 tainty: Assessing gaussian processes, deep neural nets, and deep neural operators for regression. *arXiv*  
335 *preprint arXiv:2203.04515*, 2022.
- 336 [16] A. G. Wilson and P. Izmailov. Bayesian deep learning and a probabilistic perspective of generalization.  
337 *arXiv preprint arXiv:2002.08791*, 2020.
- 338 [17] M. A. Mohamad and T. P. Sapsis. Sequential sampling strategy for extreme event statistics in nonlinear  
339 dynamical systems. *Proceedings of the National Academy of Sciences*, 115(44):11138–11143, 2018.
- 340 [18] T. P. Sapsis. Output-weighted optimal sampling for Bayesian regression and rare event statistics using  
341 few samples. *Proceedings of the Royal Society A*, 476(2234):20190834, 2020.
- 342 [19] O. Chapelle and L. Li. An empirical evaluation of Thompson sampling. *Advances in Neural Information*  
343 *Processing Systems*, 24, 2011.
- 344 [20] P. Nakkiran, G. Kaplun, Y. Bansal, T. Yang, B. Barak, and I. Sutskever. Deep double descent:  
345 Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*,  
346 2021(12):124003, 2021.
- 347 [21] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar.  
348 Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*,  
349 2020.
- 350 [22] Y. Yang, A. Blanchard, T. P. Sapsis, and P. Perdikaris. Output-weighted sampling for multi-armed  
351 bandits with extreme payoffs. *arXiv preprint arXiv:2102.10085*, 2021.
- 352 [23] L. K. Hansen and P. Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and*  
353 *Machine Intelligence*, 12(10):993–1001, 1990.
- 354 [24] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation  
355 using deep ensembles. *Advances in Neural Information Processing Systems*, 30, 2017.
- 356 [25] F. K. Gustafsson, M. Danelljan, and T. B. Schon. Evaluating scalable Bayesian deep learning methods  
357 for robust computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*  
358 *Pattern Recognition Workshops*, pages 318–319, 2020.
- 359 [26] I. Loshchilov and F. Hutter. SGDR: Stochastic gradient descent with warm restarts. *arXiv preprint*  
360 *arXiv:1608.03983*, 2016.
- 361 [27] G. Huang, Y. Li, G. Pleiss, Z. Liu, J. E. Hopcroft, and K. Q. Weinberger. Snapshot ensembles: Train  
362 1, get m for free. *arXiv preprint arXiv:1704.00109*, 2017.
- 363 [28] L. N. Smith. No more pesky learning rate guessing games. *CoRR*, *abs/1506.01186*, 5, 2015.
- 364 [29] R. B. Gramacy and H. K. H. Lee. Adaptive design and analysis of supercomputer experiments. *Tech-*  
365 *nometrics*, 51(2):130–145, 2009.
- 366 [30] A. J. Majda, D. W. McLaughlin, and E. G. Tabak. A one-dimensional model for dispersive wave  
367 turbulence. *Journal of Nonlinear Science*, 7(1):9–44, 1997.
- 368 [31] D. Cai, A. J. Majda, D. W. McLaughlin, and E. G. Tabak. Spectral bifurcations in dispersive wave  
369 turbulence. *Proceedings of the National Academy of Sciences*, 96(25):14216–14221, 1999.
- 370 [32] V. E. Zakharov, P. Guyenne, A. N. Pushkarev, and F. Dias. Wave turbulence in one-dimensional models.  
371 *Physica D: Nonlinear Phenomena*, 152:573–619, 2001.
- 372 [33] V. E. Zakharov, F. Dias, and A. Pushkarev. One-dimensional wave turbulence. *Physics Reports*,  
373 398(1):1–65, 2004.

- 374 [34] A. Pushkarev and V. E. Zakharov. Quasibreathers in the MMT model. *Physica D: Nonlinear Phenom-*  
375 *ena*, 248:55–61, 2013.
- 376 [35] W. Cousins and T. P. Sapsis. Quantification and prediction of extreme events in a one-dimensional  
377 nonlinear dispersive wave model. *Physica D: Nonlinear Phenomena*, 280:48–58, 2014.
- 378 [36] C. E. Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*,  
379 pages 63–71. Springer, 2003.
- 380 [37] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas. Taking the human out of the loop:  
381 A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- 382 [38] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, and A. Anandkumar. Neural  
383 operator: Learning maps between function spaces. *arXiv preprint arXiv:2108.08481*, 2021.
- 384 [39] L. Lu, P. Jin, and G. E. Karniadakis. DeepONet: Learning nonlinear operators for identifying differential  
385 equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*,  
386 2019.
- 387 [40] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung,  
388 R. Roscher, et al. A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342*,  
389 2021.
- 390 [41] A. F. Psaros, X. Meng, Z. Zou, L. Guo, and G. Em. Karniadakis. Uncertainty quantification in scientific  
391 machine learning: Methods, metrics, and comparisons. *arXiv preprint arXiv:2201.07766*, 2022.
- 392 [42] M. Sensoy, L. Kaplan, and M. Kandemir. Evidential deep learning to quantify classification uncertainty.  
393 *arXiv preprint arXiv:1806.01768*, 2018.
- 394 [43] A. Malinin and M. Gales. Predictive uncertainty estimation via prior networks. *arXiv preprint*  
395 *arXiv:1802.10501*, 2018.
- 396 [44] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In  
397 *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.
- 398 [45] Y. Gal and Z. Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in  
399 deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016.
- 400 [46] Y. Gal, R. Islam, and Z. Ghahramani. Deep Bayesian active learning with image data. In *International*  
401 *Conference on Machine Learning*, pages 1183–1192. PMLR, 2017.

## 402 A SIR Pandemic Model

We implement a simple Susceptible, Infected, Recovered (SIR) model,

$$\frac{dS}{dt} = -\beta(t)IS + \delta R \quad (12)$$

$$\frac{dI}{dt} = \beta(t)IS - \gamma I \quad (13)$$

$$\frac{dR}{dt} = \gamma I - \delta R, \quad (14)$$

403 where  $\delta$  is the rate of immunity loss,  $\gamma$  is the recovery rate, and  $\beta(t)$  is a stochastic infection rate. Here we  
404 take  $\delta = 0$  and  $\gamma = 0.1$  and the stochastic infection rate,  $\beta(t)$ , is defined as

$$\beta(t) = \beta_0(\theta\Phi(t) + \phi_0), \quad (15)$$

405 where  $\Phi(t)$  is found via a Karhunen-Loeve expansion of a radial basis kernel with  $\sigma_\beta^2 = 0.1$  and length scale  
406  $\ell_\beta = 1$ ,  $\beta_0 = 3 \times 10^{-9}$ , and  $\phi_0 = 2.55$ , to ensure all infection rates are non-negative. Initial conditions for  
407 the model are  $I_0 = 50$  with total population  $P = 10^8$ , and a step size of 0.1 days is used over 45 days.

## 408 B Dissipative Nonlinear Schrödinger Equation: Majda, McLaughlin, 409 and Tabak Model

410 The Majda, McLaughlin, and Tabak [30] model is a dissipative version of the nonlinear Schrödinger (NLS)  
411 equation used for studying 1D wave turbulence. It is a one-dimensional, dispersive nonlinear prototype  
412 model with intermittent events described by

$$iu_t = |\partial_x|^\alpha u + \lambda |\partial_x|^{-\beta/4} \left( \left| |\partial_x|^{-\beta/4} u \right|^2 |\partial_x|^{-\beta/4} u \right) + iDu, \quad (16)$$

413 where  $u$  is a complex scalar, exponents  $\alpha$  and  $\beta$  are chosen model parameters, and  $D$  is a selective Laplacian  
414 (described further below). This model gives rise to four-wave resonant interactions that, especially when  
415 coupled with large scale forcing and small scale damping, produces a family of spectra revealing both direct  
416 and inverse cascades [30, 31]. A realization of the MMT model is shown in figure 6 that demonstrates these  
417 complex dynamical properties. Not only does this model provide a rich dynamical response, but also presents  
418 a unique utility as a physical model for extreme ocean waves, or rogue waves [32, 33, 34]. Hence, its study is  
419 an ideal test bed for both examining the numerical difficulties of predictive models and uncovering insights  
420 to physical, real-world applications.

421 For both ease of computation and for discussion of the terms in the MMT model, we transform the  
422 equation into wavenumber space. The pseudodifferential operator  $|\partial_x|^\alpha$ , via the Fourier transform in space  
423 becomes:  $|\partial_x|^\alpha \widehat{u}(k) = |k|^\alpha \widehat{u}(k)$  where  $k$  is the wavenumber in  $x$ . This formulation may be similarly defined  
424 on a periodic domain. We choose  $\alpha = 1/2$  and  $\beta = 0$  as done in [35], reducing equation (16) to

$$\widehat{u}(k)_t = -i|k|^{1/2} \widehat{u}(k) - i\lambda |\widehat{u}(k)|^2 \widehat{u}(k) + \widehat{D}u(k) + f(k), \quad (17)$$

425 where  $f(k)$  is a forcing and  $\widehat{D}u(k)$  is a selective Laplacian of the form:

$$\widehat{D}u(k) = \begin{cases} -(|k| - k^*)^2 \widehat{u}(k) & |k| > k^* \\ 0 & |k| \leq k^* \end{cases} \quad (18)$$

426 where  $k^*$  presents the lower bound of wavenumbers subject to dissipation. For the model considered in  
427 this study we choose  $\lambda = -0.5$ ,  $k^* = 20$ ,  $f(k) = 0$ ,  $dt = 0.001$ , and a grid that is periodic between 0-1  
428 with  $N_x = 512$  grid points. To propose a stochastic and complex initial condition,  $u(x, t = 0)$ , we take the  
429 complex-valued kernel

$$k(x, x') = \sigma_u^2 e^{i(x-x')} e^{-\frac{(x-x')^2}{\ell_u}}, \quad (19)$$

430 with  $\sigma_u^2 = 1$  and  $\ell_u = 0.35$ . We then parametrize the stochastic initial conditions by a finite number of  
431 random variables,  $\boldsymbol{\theta}$ , using the Karhunen-Loeve expansion of the kernel's correlation matrix,

$$u(x, t = 0) \approx \boldsymbol{\theta} \boldsymbol{\Phi}(x), \quad \forall \quad x \in [0, 1) \quad (20)$$

432 where  $\boldsymbol{\theta} \in \mathbb{C}^m$  is a vector of complex coefficients and both the real and imaginary components of each  
433 coefficient are normally distributed with zero mean and diagonal covariance matrix  $\Lambda$ , and  $\{\Lambda, \boldsymbol{\Phi}(x)\}$  contains  
434 the first  $m$  eigenpairs of the correlation matrix. This gives the dimension of the parameter space as  $2m$  due  
435 to the complex nature of the coefficients. For all cases, the random variable  $\theta_i$  is restricted to a domain  
436 ranging from -6 to 6, in that 6 standard deviations in each direction from the mean.

## 437 C Gaussian Process Regression

438 For low-dimensional problems, Gaussian process (GP) regression [36] is the “gold standard” for Bayesian  
439 design. There are several attractive qualities of GPs. They are agnostic to the details of a black-box process  
440 (just like neural networks described next) and they clearly quantify both the uncertainty in the model and  
441 the uncertainty associated with noise. GPs are also relatively easy to implement and cheap to train.

442 A Gaussian process  $\bar{f}(\boldsymbol{\theta})$ , where  $\boldsymbol{\theta}$  is a random variable, is completely specified by its mean function  
443  $m(\boldsymbol{\theta})$  and covariance function  $k(\boldsymbol{\theta}, \boldsymbol{\theta}')$ . For a dataset  $\mathcal{D}$  of input-output pairs ( $\{\boldsymbol{\Theta}, \mathbf{y}\}$ ) and a Gaussian

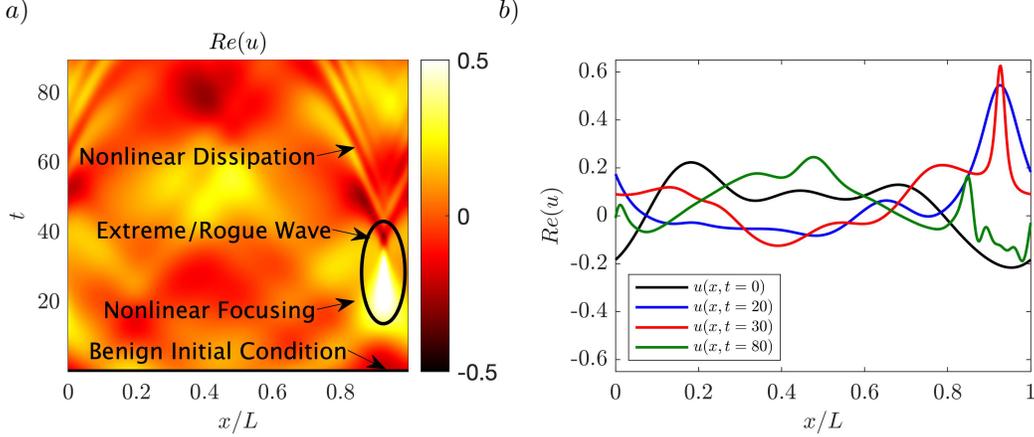


Figure 6: **The NLS emits many nonlinear phenomena in finite time, from focusing to dissipation.** *a)* A realization of the NLS model ( $\alpha = 1/2, \beta = 0, \lambda = -0.5$ ) we discover with DNO+BED, where a benign initial condition leads to an extreme, rogue wave. *b)* Wave height of selected times from *a)*.

444 process with constant mean  $m_0$ , the random process  $\bar{f}(\theta)$  conditioned on  $\mathcal{D}$  follows a normal distribution  
 445 with posterior mean and variance

$$\mu(\theta) = m_0 + k(\theta, \Theta) \mathbf{K}^{-1} (\mathbf{y} - m_0) \quad (21)$$

446

$$\sigma^2(\theta) = k(\theta, \theta) - k(\theta, \Theta) \mathbf{K}^{-1} k(\Theta, \theta) \quad (22)$$

447 respectively, where  $\mathbf{K} = k(\Theta, \Theta) + \sigma_\epsilon^2 \mathbf{I}$ . Equation (21) can predict the value of the surrogate model at any  
 448 point  $\theta$ , and (22) to quantify uncertainty in prediction at that point [36]. Here, we chose the radial-basis-  
 449 function (RBF) kernel with automatic relevance determination (ARD),

$$k(\theta, \theta') = \sigma_f^2 \exp \left[ -(\theta - \theta')^\top \mathbf{L}^{-1} (\theta - \theta') / 2 \right], \quad (23)$$

450 where  $\mathbf{L}$  is a diagonal matrix containing the lengthscales for each dimension and the GP hyperparameters  
 451 appearing in the covariance function ( $\sigma_f^2$  and  $\mathbf{L}$  in (22) are trained by maximum likelihood estimation).

452 One setback from the above expression is the inference step in GP regression, where each iteration requires  
 453 the inversion of the matrix  $\mathbf{K}$ . Typically performed by Cholesky decomposition, the inversion cost scales  
 454 as  $O(n^3)$ , with  $n$  being the number of observations [36, 37]. This means that as problems grow to infinite  
 455 dimensions, inevitably requiring larger datasets, GPs become prohibitively costly. We will show here that  
 456 as the need for data increases, GPs become much more computationally intensive than our next surrogate  
 457 model, neural networks.

## 458 D Deep Neural Operators and DeepONet

459 Unlike GPs, Deep Neural Operators or Deep Neural Networks, do not suffer from data-scaling challenges  
 460 and are our primary model class for consideration in this manuscript. Deep Neural Networks, when cast as  
 461 neural operators[38], are specifically well-suited for characterizing infinite dimensional systems, as they may  
 462 map functional inputs to functional outputs. Although our work is general for any neural network approach,  
 463 we leverage the architecture proposed by [7] for approximating nonlinear operators: *DeepONet*.

464 DeepONet seeks approximations of nonlinear operators by constructing two deep neural networks, one  
 465 representing the input function at a fixed number of sensors and another for encoding the “locations” of  
 466 evaluation of the output function. The first neural network, termed the “branch”, takes input functions,  
 467  $u$ , observed at discrete sensors,  $x_i, i = 1 \dots m$ . These input functions can take on several representations,  
 468 such as initial conditions (i.e.  $u_0$ ) or forcing functions. The second neural network, termed the “trunk”,  
 469 should be seen as an encoder for inherent qualities of the operator, denoted as  $z$  (referred to as  $y$  in [7], but

470 changed here for typical active sampling notation). For example, a variable coefficient or exponent in the  
 471 true nonlinear operator,  $G$ , alternatively, and the usual use case for DeepONet, the trunk variable can refer  
 472 to the evaluation of the operation to an arbitrary point in time and/or space. Together, these networks seek  
 473 to approximate the nonlinear operation upon  $u$  and  $z$  as  $G(u)(z) = y$ , where  $y$  will denote the scalar output  
 474 from the  $u, z$  input pair.

475 Given a set of input-output pairs,  $\{[\mathbf{u}, \mathbf{z}], G(\mathbf{u})(\mathbf{z})\}$ , DeepONet seeks to minimize the difference between  
 476 the true operator,  $G(u)(z)$ , and the dot product between two neural networks  $\mathbf{g}(u)$  and  $\mathbf{f}(z)$ . These network’s  
 477 level of expressivity is governed by the number of neurons ( $n$ ), the number of layers ( $l_b, l_t$  for branch and  
 478 trunk, respectively), and activation functions. Under sufficient training DeepONet can meet any arbitrary  
 479 error,  $\epsilon$ , as

$$|G(u)(z) - \underbrace{\langle \mathbf{g}(u(x_1), u(x_2), \dots, u(x_m)), \mathbf{f}(z) \rangle}_{\text{branch trunk}}| < \epsilon. \quad (24)$$

480 Thus, DeepONet can, to an arbitrarily small precision, approximate an infinite dimensional nonlinear oper-  
 481 ator  $G(u)(z)$ . However, the functions  $\mathbf{g}$  and  $\mathbf{f}$  are typically trained under the assumption of plentiful data.  
 482 We are interested in how DeepONet can be optimally trained, in that with the least amount of data, for  
 483 discovering and quantifying extreme events.

484 DeepONet is attractive for these tasks as neural nets generalize well and are incredibly fast to evaluate  
 485 (compared to their experimental/simulation counterparts) for arbitrarily chosen points. However, there is  
 486 little work on how one optimally selects the best samples to train a neural net. Previously, the approach taken  
 487 has included creating an appropriate basis for specific operators that are used to train DeepONet [39]. For  
 488 the problems we are interested (i.e. rare and extreme events), this appropriate basis is unknown. Therefore,  
 489 we envision DeepONet will provide a flexible model to learn seen data, and then provide early predictions  
 490 of where danger and uncertainty lie in the input/output space through DeepONet’s parameterization, that  
 491 we may then query the underlying system to best inform DeepONet.

## 492 D.1 Approaches for quantifying uncertainty in Deep Neural Networks

493 There are several techniques for quantifying uncertainty in neural networks and three categorizations of these  
 494 techniques we wish to mention: single deterministic networks, Bayesian neural network (BNN) inference ap-  
 495 proaches, and ensemble methods (see [40] and [41] for comprehensive reviews of these methods for uncertainty  
 496 quantification). Unfortunately, a complex array of advantages and disadvantages of each approach provides  
 497 no clear favorite. Single deterministic methods [42, 43] use one forward pass of a deterministic network  
 498 to learn both the mean and variance of a labeled output. Using only one model leads to cheap training  
 499 and evaluation. However, this single opinion of the underlying system results in substantial sensitivity, an  
 500 unattractive quality for regression problems quantifying physical instabilities/extreme events. Bayesian neu-  
 501 ral networks [44, 45, 46, 16] encompass a broad variety of stochastic DNNs that combine Bayesian inference  
 502 theory with the expressiveness, scalability, and predictive performance enjoyed by deep neural networks. Al-  
 503 though the supporting theory behind Bayesian methods, as well as empirical results, imparts faith that such  
 504 models will lead to the greatest chance of success, they do come with a significant disadvantage, complexity.  
 505 BNNs are significantly more complex than standard neural networks and can be exhausting to train, making  
 506 them difficult to implement. From an academic viewpoint, this can be overcome. However, as a study  
 507 concerned with translating DNNs and active learning to practical engineering systems that solve real-world  
 508 problems, we opt for using the straight-forward ensemble approach. An approach that seats itself between  
 509 the single deterministic model and the infinite model representations of BNNs.

## 510 E DeepONet Setup

511 Application of the NLS and SIR models to the DeepONet architecture requires a well-posed distinction  
 512 between the functions and parameters that belong in the branch and trunk networks. Figure 7 provides a  
 513 visual of the function/parameter delineations and the architecture for each modeling task.

514 In this work, we provide all input initial conditions as functions to the branch network. Because of the  
 515 requirement that all inputs to DeepONet take on real values, the NLS setup requires that we split the input  
 516 function,  $u$ , into its real ( $u_r$ ) and imaginary ( $u_i$ ) components. It is critical to keep both components, as

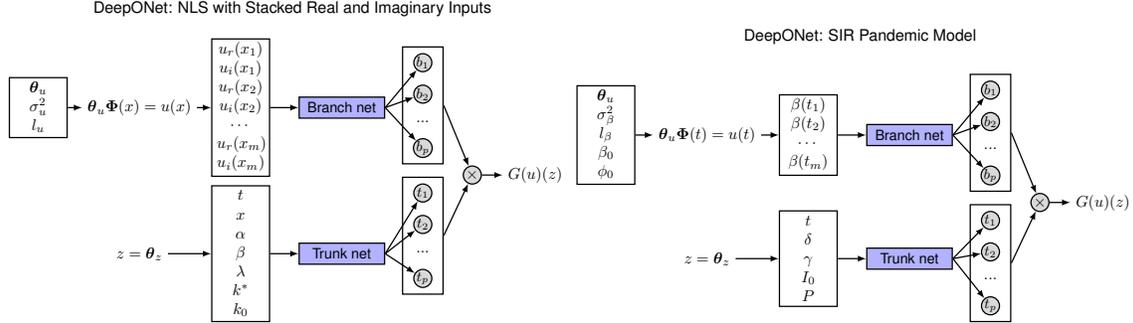


Figure 7: The application of the NLS and SIR operators to DeepONet.

Case	Neurons	Branch Layers	Ensemble Size	Experiments
$2 - 6D$	200	5	10	10
$8D$ Batching	200	5	10	10
$8D$ Ensemble, $n \leq 2500$	200	5	2,4,8,16	10
$8D$ Ensemble, $n > 2500$	200	6	2,4,8,16	10
$20D$ , $n \leq 2500$	200	5	2	25
$20D$ , $n > 2500$	200	6	2	25

Table 1: Hyperparameters used for all cases. All trunk layers = 1 due to their redundancy.

each contains unique and coupled information that is propagated in the NLS operation (this is not necessary for the real-valued SIR model). We then stack these components as a vector by  $x$  position/sensor with the imaginary component directly following the real component at each  $x$  position. Technically, the ordering of the inputs does not matter, due to the linear nature of the first layer, but the ordering of the inputs must remain consistent. For increased speed, we also reduce the sensor number of  $x_m$  sensors from the 512 used for the direct NLS calculation to 128 points for DeepONet (this corresponds to 256 total inputs values because of the real and imaginary components) and 125 points for the SIR infection rate. We also give the length scale and variance as input parameters, as they define the kernel that emits the Karhunen-Loeve expansion functions. The input functions directly recognize adjustments to these parameters.

The trunk contains parameters that are intrinsic to the NLS and SIR operations. For NLS, these comprise both time and space  $(t, x)$  and the chosen constants  $\alpha, \beta, \lambda, k^*$ . Regarding BED, we may assign each of these parameters to an appropriate prior distribution, such as uniform for space and time. In this study, we choose all values to remain constant and although the trunk is straightforward to implement; it is simply a redundant network here.

We provide the hyperparameters and other quantities used for the various NLS cases performed throughout the manuscript in table 1. In addition to the hyperparameters, we use the ReLu activation function, a learning rate of  $l_r = 0.001$ , and 1000 epochs for each training procedure. For the SIR search, the same hyperparameters are used with 6 branch layers, 1 redundant trunk layer, and 200 neurons.

## F Monte Carlo Minimization at $20D$

Here we show that acquisition points found through Monte Carlo (MC) searches are consistently superior than those selected by minimizers using gradient descent algorithms (when reasonably similar computation times are considered). To demonstrate this behavior, we test four methods for finding 50 batched acquisition points for the  $20D$  NLS case at 500, 2500, and 5000 training samples,  $N = 2$  ensemble members, and over 25 independent experiments. The four approaches are listed below, with each case choosing the 50 best points with algorithm 2:

1. L-BFGS-B algorithm, implemented within the `scipy` Python package,  $10^2$  random LHS initial points.
2. MC with  $10^5$  random uniform initial points (sampled from a uniform distribution for speed)

Samples	L-BFGS-B	MC: $10^5$	MC+L-BFGS-B	MC: $10^6$
500	$122 \pm 7$	$3.6 \pm 0.1$	$123 \pm 5$	<b><math>27 \pm 1</math></b>
2500	$126 \pm 6$	$4.0 \pm 0.2$	$125 \pm 5$	<b><math>26 \pm 2</math></b>
5000	$127 \pm 8$	$3.7 \pm 0.1$	$135 \pm 7$	<b><math>27 \pm 3</math></b>

Table 2: **MC is faster than built-in optimizers.** Mean compute times ( $n = 25$ ) for batch minimization ( $n_b = 50$ ) of the the 20D US-LW acquisition function with  $\pm$  one standard deviation. Best values are bold.

544 3. MC with  $10^5$ , best 100 points (algorithm 2) passed to L-BFGS-B.

545 4. MC with  $10^6$  random uniform initial points.

546 Table 2 and table 3 provide the mean computation times and scores for all cases, respectively. The mean  
547 computation times ( $\pm$  one standard deviation) are clearly faster for MC methods, by 5-fold, compared to  
548 any approach with the L-BFGS-B algorithms. We also note that only the first ensemble member is used for  
549 computing and querying  $p_\mu$  for speed, as well as likely better performance.

550 Table 3 provides the mean scores (over all chosen batch points and experiments,  $n = 1250$ ), as well  
551 as the mean best and worst point for each experiment ( $n = 25$ ). It is critical to note that at these high  
552 dimensions the difference between acquisition scores can be *hundreds* of orders of magnitude different (this  
553 is a direct result of weights define by the ratio of two PDFs computed over 20 dimensions). To navigate this  
554 computational challenge, we take the  $\log_{10}$  of the scores and apply a negative sign such that lower scores  
555 are optimal (i.e.  $-\log_{10}(a(\boldsymbol{\theta}))$ ). Considering the mean of all chosen points, we see that a naive use of the  
556 L-BFGS-B minimizer finds scores that are  $\approx 50$  orders of magnitude worse than the  $10^6$  MC approach,  
557 regardless of the training complexity of the model. The L-BFGS-B minimizer only provides a marginal  
558 advantage for finding the best point in any random instance, however, the standard deviation is 30 orders  
559 of magnitude and suggests the approach is terribly inconsistent. Juxtaposing the min and max results, we  
560 see that the MC method only varies by approximately 10 orders of magnitude compared to the 80 orders of  
561 the L-BFGS-B.

562 Overall, the MC method provides greater consistency in finding several attractive acquisition points,  
563 while also enjoying ease of implementation and better computational efficiency. Of course, if given enough  
564 computational resources and time, the L-BFGS-B method permits pinpointing exceptional acquisition points.  
565 Approaches combining both MC and L-BFGS-B, where the MC points are further optimized (as also imple-  
566 mented here) or performed from more LHS points, can either refine or enrich the set of acquisition points.

567 We believe the results found here are chiefly because of the non-convexity of the acquisition function.  
568 We may recall the highly non-convex behavior of the 2D acquisition fields in figure 2b), even with as little as  
569  $\approx 10$  samples. This non-convex nature emits many local minima that requires many initial points to provide  
570 confidence that the chosen optima is nearly global. This makes the task for built in minimizers extremely  
571 difficult, especially for high dimensional problems, and lead to minimizations that become easily fooled or  
572 stuck. This difficulty is so challenging that the sparse MC sampling of only  $10^6$  points, an extremely sparse  
573 sampling of a 20D space, easily outperforms the minimizers.

<b>Samples</b>	L-BFGS-B	MC: $10^5$	MC+L-BFGS-B	MC: $10^6$
Mean ( $n = 1250$ )				
500	$110 \pm 18$	$70 \pm 5$	$66 \pm 6$	<b><math>60 \pm 4</math></b>
2500	$104 \pm 16$	$67 \pm 5$	$63 \pm 5$	<b><math>56 \pm 4</math></b>
5000	$110 \pm 17$	$76 \pm 4$	$67 \pm 8$	<b><math>64 \pm 4</math></b>
Min $n = 25$				
500	<b><math>42 \pm 32</math></b>	$57 \pm 6$	$53 \pm 16$	$52 \pm 5$
2500	<b><math>44 \pm 32</math></b>	$53 \pm 4$	$48 \pm 10$	$47 \pm 6$
5000	$53 \pm 21$	$60 \pm 5$	<b><math>48 \pm 11</math></b>	$54 \pm 3$
Max $n = 25$				
500	$126 \pm 7$	$73 \pm 3$	$70 \pm 3$	<b><math>63 \pm 3</math></b>
2500	$117 \pm 6$	$70 \pm 3$	$65 \pm 3$	<b><math>60 \pm 3</math></b>
5000	$125 \pm 5$	$80 \pm 2$	$72 \pm 2$	<b><math>67 \pm 2</math></b>

Table 3: **MC optimization provides consistently superior acquisition scores over built-in optimizers.** Top rows: The mean and standard deviation of all 1250 chosen acquisition points for the four methods on DNOs trained with 500, 2500, and 5000 samples. Middle rows: The mean and standard deviation of the best acquisition points for the 25 experiments. Bottom rows: The mean and standard deviation of the worst chosen acquisition points (i.e. the least optimal point of the batch). All values are  $-\log_{10}(a(\theta))$ . Best values are bold.

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [EPickeringEPCflat.pdf](#)