

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Halonen, Vilho; Pölönen, Ilkka

Title: Quantification of Errors Generated by Uncertain Data in a Linear Boundary Value Problem Using Neural Networks

Year: 2023

Version: Accepted version (Final draft)

Copyright: © 2023 the Authors

Rights: CC BY 4.0

Rights url: <https://creativecommons.org/licenses/by/4.0/>

Please cite the original version:

Halonen, V., & Pölönen, I. (2023). Quantification of Errors Generated by Uncertain Data in a Linear Boundary Value Problem Using Neural Networks. *SIAM/ASA Journal on Uncertainty Quantification*, 11(4), 1258-1277. <https://doi.org/10.1137/22M1538855>

Quantification of Errors Generated By Uncertain Data in a Linear Boundary Value Problem Using Neural Networks*

Vilho Halonen[†] and Ilkka Pölönen[†]

Abstract. Quantifying errors caused by indeterminacy in data is currently computationally expensive even in relatively simple PDE problems. Efficient methods could prove very useful in for example scientific experiments done with simulations. In this paper, we create and test neural networks which quantify uncertainty errors in the case of a linear one-dimensional boundary value problem. Training and testing data is generated numerically. We created three training datasets, three testing datasets and trained four neural networks with differing architectures. The performance of the neural networks is compared to known analytical bounds of errors caused by uncertain data. We find that the trained neural networks accurately approximate the exact error quantity in almost all cases and the neural network outputs are always between the analytical upper and lower bounds. The results of this paper show that after a suitable dataset is used for training even a relatively compact neural network can successfully predict quantitative effects generated by uncertain data. If these methods can be extended to more difficult PDE problems they could potentially have a multitude of real-world applications.

Key words. PDE, Uncertainty Quantification, Machine Learning, Neural Network

MSC codes. 65G99, 68T01, 65L10

1. Introduction. Mathematical models in real-world applications are always burdened by uncertainty in measurements. Solutions to real-world problems are usually created by numerical methods such as the finite element method [2] and the data used is never exactly known. Uncertainty in data causes errors in the outputs of mathematical models. Currently in PDE problems the methods used to quantify these errors include probabilistic methods [10], the worst case scenario method [5] and a posteriori error estimates [6].

There are two major issues with current methods. Analytical methods are not always available and even if they exist the results may be too coarse. Numerical methods can suffer from unreasonably high computational costs. If a solution seems good enough an engineer may disregard uncertainty error estimation completely.

The problem of uncertainty analysis is close to the area of scientific research where modern machine learning models such as neural networks are used. In this paper, we investigate the possibility of using neural networks to quantify errors generated by uncertain data. First we create datasets which have as inputs the uncertainty in data and as outputs the uncertainty error represented by some error metric. These datasets are used to train various different neural networks. To analyse effects of uncertain data we chose an example differential equation which is presented in section 3. For the chosen problem, analytical error bounds were derived in [4].

The uncertainty information is not a discrete set of values which can be directly input to a neural network so we need to characterize the uncertainty in an approximated way. This is

*Submitted to the editors 07.12.2022.

[†]University of Jyväskylä, Faculty of Information Technology. Mattilanniemi 2, 40100, Jyväskylä, Finland.

40 done by scanning the set of possible values for parameters on some amount of specific points.
 41 The prototype problem we use has uncertain parameter functions which are one-dimensional,
 42 so the scanning is done on the edge of the maximum and minimum values of a function in the
 43 indeterminacy set. In our case, we used eight points for this scanning operation (see Figure
 44 2).

45 The paper is organized as follows. First we introduce the uncertainty problem in an
 46 abstract setting in section 2. In section 3 we show the prototype problem used, methods
 47 to create datasets, the neural network models and how network performance is evaluated.
 48 Results are given in section 4, and the conclusions follow in section 6.

49 **2. Problem Setting.** Consider an abstract differential problem

$$50 \quad (2.1) \quad \mathcal{A}u = f,$$

52 where $u \in V$ is the unknown function in some suitable solution space, \mathcal{A} is a differential
 53 operator and f is the source term. In practice, the differential operator \mathcal{A} and source term
 54 f depend on measured data which is subject to inaccuracies in measurements. For example,
 55 if 2.1 were a linear elasticity problem the operator \mathcal{A} would contain information about the
 56 Lamé parameters of the deformed material and f would describe the outside forces subjected
 57 to the material. Such measured quantities are only known up to some accuracy. A solution
 58 to 2.1 is therefore not unique but rather there is a set of possible solutions.

59 Assume that instead of knowing \mathcal{A} and f exactly we only know that

$$60 \quad (2.2) \quad (\mathcal{A}, f) \in \mathcal{D} := \mathcal{D}_{\mathcal{A}} \times \mathcal{D}_f.$$

62 We call the set \mathcal{D} the *set of admissible data*. Each element of the set \mathcal{D} may produce a different
 63 solution u to the problem 2.1. For this purpose we define the set of solutions which is the
 64 image of the following solution mapping.

65 **Definition 2.1.** *The solution mapping $\mathcal{S} : \mathcal{D} \rightarrow V$ is a function such that $\mathcal{S}(\mathcal{A}, f) = u$,*
 66 *where u is the exact solution to 2.1. The image $\mathcal{S}(\mathcal{D})$ is called the solution set.*

67 If possible, we would like to know quantitatively what the solution set $\mathcal{S}(\mathcal{D})$ is like. For this
 68 purpose the quantity of interest is the diameter of this set with respect to some suitable norm.

Definition 2.2. *The diameter of the solution set is*

$$Diam(\mathcal{S}(\mathcal{D})) := \sup_{u_1, u_2 \in \mathcal{S}(\mathcal{D})} \|u_1 - u_2\|,$$

69 where the norm $\|\cdot\|$ is a suitable norm in the vector space V .

70 The quantity $Diam(\mathcal{S}(\mathcal{D}))$ is useful for two things. Firstly, if there is a priori knowledge of what
 71 this quantity is expected to be, the accuracy of measurements related to \mathcal{D} can be considered
 72 suitable or unsuitable. This quantity can also give us an accuracy limit on approximation
 73 errors. If approximation errors are lower than the diameter the approximation may already
 74 be within the solution set so sharpening approximation schemes is a wasted effort (see Figure
 75 1). Our research objective is to create neural networks which use the indeterminacy set \mathcal{D} as
 76 input and outputs an approximation of the diameter $Diam(\mathcal{S}(\mathcal{D}))$.




Figure 1. Since approximation error is less than the diameter sharpening the approximation scheme is not useful.

77 **3. Methods.** We investigate the effects of uncertain data in the context of a prototype
 78 problem which is presented in section 3.1. Our process of creating neural network models
 79 consists of three distinct steps. First, we describe and implement creation of training and
 80 test data in section 3.2. Afterwards we create and train neural networks on the made data in
 81 section 3.4. In section 3.5, we explain how network performance is tested.

82 **3.1. Prototype problem.** The effects of indeterminacy are studied in the case of the
 83 following boundary value problem. Problem \mathcal{P} : Find $u \in H_0^1(0, 1)$, such that

$$84 \quad (3.1) \quad (\alpha(x)u'(x))' - \beta(x)u(x) = f(x),$$

$$85 \quad (3.2) \quad u(0) = u(1) = 0,$$

$$86 \quad (3.3) \quad 0 < \alpha_{\ominus} < \alpha(x) < \alpha_{\oplus}, \quad 0 < \beta_{\ominus} < \beta(x) < \beta_{\oplus}, \quad f \in L^2(0, 1).$$

88 Uncertainty in (3.1)-(3.3) is introduced by assuming that α , β and f are not known exactly.
 89 Instead we only know that

$$90 \quad (\alpha, \beta, f) \in \mathcal{D} = \mathcal{D}_{\alpha} \times \mathcal{D}_{\beta} \times \mathcal{D}_f,$$

92 where

$$93 \quad \mathcal{D}_{\alpha} := \{\alpha_{\circ} + \delta_1 g : \|g\|_{\infty} \leq 1\},$$

$$94 \quad \mathcal{D}_{\beta} := \{\beta_{\circ} + \delta_2 g : \|g\|_{\infty} \leq 1\},$$

$$95 \quad \mathcal{D}_f := \{f_{\circ} + \delta_3 g : \|g\|_{\infty} \leq 1\}.$$

97 The functions α_{\circ} , β_{\circ} and f_{\circ} are referred to as the "mean" or "central" elements of the
 98 admissible data and the functions δ_1 , δ_2 and δ_3 are the maximum perturbations from them.
 99 The image of the solution mapping (see 2.1) is in the space $H_0^1(0, 1)$, so $\mathcal{S} : \mathcal{D} \rightarrow H_0^1(0, 1)$. In
 100 this space the suitable norm that we use for the diameter $\text{Diam}(\mathcal{S}(\mathcal{D}))$ (see 2.2) is the energy
 101 norm

$$102 \quad (3.4) \quad \|u\|_{\circ} := \left(\int_0^1 \alpha_{\circ} |u'|^2 + \beta_{\circ} |u|^2 dx \right)^{\frac{1}{2}},$$

104 where $u_{\circ} = \mathcal{S}(\alpha_{\circ}, \beta_{\circ}, f_{\circ})$.

105 For this one-dimensional problem the input for a neural network is formed by scanning
 106 each of the sets \mathcal{D}_{α} , \mathcal{D}_{β} and \mathcal{D}_f on the edge of what the maximum and minimum function in
 107 these sets are (see Figure 2) on eight equidistant points on the domain $[0, 1]$.




Figure 2. Example of how the set \mathcal{D}_α is turned into a discrete input for a neural network. The y-coordinate of each red dot is put into a vector of length 16 in a chosen order. This vector is part of the neural network input together with similar vectors scanned from the sets \mathcal{D}_β and \mathcal{D}_f .

108 **3.2. Creating Data.** We created three training datasets and three testing datasets. The
 109 training datasets consist of examples where the admissible datasets $\mathcal{D}_\alpha, \mathcal{D}_\beta$ and \mathcal{D}_f have a
 110 piecewise constant maximum and minimum element (see Figure 4) Examples in the made
 111 test datasets are less restricted in terms of smoothness and oscillation. Rather than being
 112 piecewise constant on their edges the edge is smooth and can oscillate (see Figure 3) more and
 113 more for the subsequent test datasets. This will let us see how the networks generalize when
 114 trained with a fairly constrained dataset. It is also noticeably easier and faster to generate
 115 training examples with tighter constraints. In our case, the piecewise constraints make it
 116 so that different solutions in the solution set can be found by solving relatively small linear
 117 systems of equations which is done very quickly on modern computers.

118 First the method of generating a single example is discussed and in the following subsec-
 119 tions the specific characteristics of each dataset are presented. A training example consists of
 120 the input which is a discrete version of the admissible dataset \mathcal{D} and the output A which is
 121 an approximation of the quantity $\text{Diam}(\mathcal{S}(\mathcal{D}))$. Discrete inputs are created by scanning the
 122 admissible dataset on eight points on the domain (see Figure 2).

123 To approximate the diameter we randomly select a finite amount of elements from the
 124 admissible dataset and compute the solution for each choice. From the resulting finite solution
 125 set the diameter can be approximated by brute force calculation of the maximum energy
 126 distance between two elements. In our case, the finite set was chosen to have 2000 solutions.
 127 In testing different sizes it was found that after around size 1500 the diameter approximation

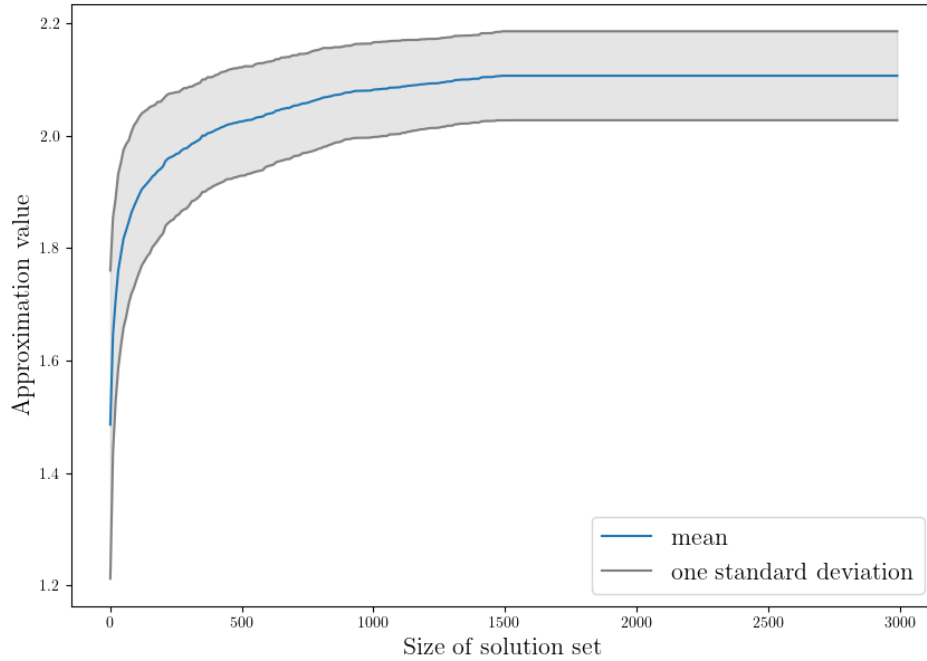


Figure 3. Diameter approximation for solution sets of different sizes. The approximation was computed 200 times for the different sizes of solution sets [10, 20, \dots , 3000]. After 1500 solutions are reached the approximation did not change so when making training examples we use 2000 solutions to approximate the diameter.

128 tends to not increase (see Figure 3). To be safer we go a little beyond and choose 2000.
 129 Note however that if we are unlucky even with such an amount of solutions the diameter
 130 approximation may be inaccurate. For our purposes it is sufficiently accurate.

131 It is not trivial to decide what the random elements chosen from \mathcal{D} should be. If they are
 132 too complex generating solutions may be slow. If they are too simple the finite set will not be
 133 a good representation of \mathcal{D} . For the training datasets we chose piecewise constant functions.
 134 For such functions the solutions can be generated very fast. The specifics of how the problem
 135 (3.1)-(3.3) becomes a linear system when using piecewise constant functions α, β and f are
 136 described in 3.2.1. The reason we chose to use these piecewise datasets for training is to
 137 see if a somewhat limited type of training data can train a network such that it generalizes
 138 to more complex functions. If the PDE at hand were more complicated it may provide a
 139 substantial increase in computational cost to make a training dataset that is somehow limited
 140 in complexity.

141 After the finite solution set has been generated the diameter is calculated by brute force
 142 computation which compares each pair of solutions in the solution set and computes the
 143 energy distance between them. The maximum distance found is the approximated diameter
 144 which is used as the output of the training example. This procedure is rather computationally
 145 expensive and takes almost all of the time used to generate a single example.

146 **3.2.1. Training Datasets.** In all of the datapoints in the training datasets the mean
147 functions α_o , β_o , f_o and the indeterminacy functions δ_i are piecewise constant such that

$$\begin{aligned} 148 \quad (3.5) \quad & \alpha_o(x) = \alpha_k, \text{ when } x \in [x_{k-1}, x_k], \\ 149 & \beta_o(x) = \beta_k, \text{ when } x \in [x_{k-1}, x_k], \\ 150 & f_o(x) = f_k, \text{ when } x \in [x_{k-1}, x_k], \\ 151 & \delta_i(x) = \delta_i^k, \text{ when } x \in [x_{k-1}, x_k], \quad i = 1, 2, 3. \end{aligned}$$

153 The amount of subintervals is the same for the first two datasets and increased for the last
154 dataset. The indeterminacy functions have a magnitude limit given by

$$155 \quad (3.6) \quad \delta_1(x) \leq 0.3\alpha_o(x), \quad \delta_2(x) \leq 0.3\beta_o(x), \quad \delta_3(x) \leq 0.3f_o(x).$$

157 From the admissible dataset with the above constraints we pick 2000 individual solutions
158 to (3.1)-(3.3) by sampling random functions α , β and f from the admissible dataset. In every
159 sample the functions α , β and f are also chosen to be piecewise constant. This lets us generate
160 each solution in this set of 2000 solutions by solving linear system of equations. How this is
161 done is described as follows:

162 Assume that the functions α , β , f are constant on the subintervals (x_i, x_{i+1}) , $i = 0, N$.

$$\begin{aligned} 163 & \alpha(x) = \alpha_k, \text{ when } x \in [x_{k-1}, x_k], \\ 164 & \beta(x) = \beta_k, \text{ when } x \in [x_{k-1}, x_k], \\ 165 & f(x) = f_k, \text{ when } x \in [x_{k-1}, x_k]. \end{aligned}$$

167 A differential equation of the form

$$168 \quad \alpha u'' - \beta u = f$$

170 where α , β and f are constant has a general solution of the form

$$171 \quad (3.7) \quad u = c_1 e^{(\sqrt{\beta/\alpha})x} + c_2 e^{-(\sqrt{\beta/\alpha})x} - \frac{f}{\beta}.$$

173 With n subintervals we will have n equations of the form (3.7). We denote the solutions
174 u_1, u_2, \dots, u_n such that

$$175 \quad u_k(x) = c_{k1} e^{(\sqrt{\beta_k/\alpha_k})x} + c_{k2} e^{-(\sqrt{\beta_k/\alpha_k})x} - \frac{f_k}{\beta_k}.$$

177 This gives us $2n$ unknown parameters c_{ij} . On the boundary of each subinterval the value
178 of the solution u and the value of the flux $\alpha u'$ must coincide. On top of that we have the
179 boundary conditions $u(a) = A$, $u(b) = B$ at the boundary of the full interval. This gives us
180 the required $2n$ equations to solve this system. This is a linear system with the matrix form
181 $Ac = F$. Solving this system gives us the unknowns c_{ij} . For this reason it is very efficient to
182 create a representation of the solution cloud by sampling piecewise constant functions from
183 the admissible data \mathcal{D} .




Figure 4. Example of what admissible sets \mathcal{D}_α in the training datasets can look like. The constraint on each subinterval is that the functions in the admissible set are between two constants and the amount of subintervals is chosen for each training dataset. The y-coordinates of the red dots are used as inputs for neural networks.

184 We created three training datasets and the examples in each dataset have the following
 185 constraints on top of the indeterminacy magnitude constraint 3.6 and the piecewise constant
 186 mean and indeterminacy function constraint 3.2.1.

187 Training dataset 1:

- 188 • The domain $[0, 1]$ is split into four pieces:
- 189 $[x_0, x_1, x_2, x_3, x_4] = [0, 0.25, 0.5, 0.75, 1]$.
- 190 • $4 \leq \alpha_o(x) \leq 8, 4 \leq \beta_o(x) \leq 8, 0 \leq f_o(x) \leq 10$.

191 Training dataset 2:

- 192 • The domain $[0, 1]$ is split into four pieces:
- 193 $[x_0, x_1, x_2, x_3, x_4] = [0, 0.25, 0.5, 0.75, 1]$.
- 194 • $4 \leq \alpha_o(x) \leq 14, 4 \leq \beta_o(x) \leq 14, -20 \leq f_o(x) \leq 20$.

195 Training dataset 3:

- 196 • The domain $[0, 1]$ is split into eight pieces:
- 197 $[x_0, x_1, x_2, \dots, x_8] = [0, 0.125, 0.25, \dots, 1]$.
- 198 • $4 \leq \alpha_o(x) \leq 14, 4 \leq \beta_o(x) \leq 14, -20 \leq f_o(x) \leq 20$.

199 Each training dataset contains a total of 2000 datapoints. A single datapoint consists of the
 200 discrete scanned vector of length 48 generated by scanning the admissible set \mathcal{D} (see Figure
 201 4) and the approximated diameter of the solution set.

202 **3.2.2. Test Datasets.** We created three test datasets. The difference compared to the
 203 training datasets is the smoothness of the admissible datasets \mathcal{D} , the amount of oscillation in

204 the domain and the range in which the mean functions α_o, β_o and f_o are. The first two test
 205 sets have the mean functions in the same ranges as all of the training datasets while the last
 206 test set has mean functions which have higher values than in any of the training examples.
 207 Oscillation is included in all three being the lower in the first and third set and substantially
 208 higher in the second test set. With these choices of test sets we can see if the piecewise
 209 constrained training data is sufficient to generalize to less constrained admissible data \mathcal{D} .

210 The admissible datasets for test sets are generated by randomly picking some number of
 211 points between a maximum and minimum value and fitting a spline curve to these points.
 212 This spline curve is used as the mean function of the admissible dataset. The indeterminacy
 213 functions δ_i are created in the same way. This produces smoother admissible datasets as
 214 shown in figure 5. All test sets have 130 datapoints.

215 The values which we can pick to create different kinds of test sets are:

- 216 1. The range of values for $\alpha_o(x) \in [\alpha_-, \alpha_+]$ (similarly for β_o and f_o).
- 217 2. The range of values for $\delta_i \in [\delta_{i-}, \delta_{i+}]$.
- 218 3. The amount of random points used to make the splines (controls how much oscillation
 219 is possible).

220 These values are picked as follows for each set:

221 Test set 1:

- 222 • $\alpha_o, \beta_o \in [4, 8]$
- 223 • $f_o \in [4, 10]$
- 224 • $\delta_i \in [0.1, 0.3]$
- 225 • Amount of spline points is 4

226 Test set 2:

- 227 • $\alpha_o, \beta_o \in [4, 8]$
- 228 • $f_o \in [4, 10]$
- 229 • $\delta_i \in [0.1, 0.3]$
- 230 • Amount of spline points is 9

231 Test set 3:

- 232 • $\alpha_o, \beta_o \in [14, 24]$
- 233 • $f_o \in [20, 40]$
- 234 • $\delta_i \in [0.1, 0.3]$
- 235 • Amount of spline points is 4

3.3. Neural Networks. A feed-forward neural network is a machine learning model which can be used for a variety of classification and regression problems. The structure of such a model has an input layer, one or more hidden layers and an output layer. The input layer consists of as many neurons as there are dimensions in a datapoint. Hidden layers consist of a chosen number of neurons on each layer and the output layer in the case of a regression problem has a single neuron from which the approximated output value is extracted. In a dense feed-forward neural network like the ones we used each layer is "fully-connected" to the next layer such that each neuron is connected to every neuron of the next layer and adjusted by a weight parameter w . Each neuron in the hidden layers performs the operation

$$n(\mathbf{x}) = a(\mathbf{w}^T + b),$$

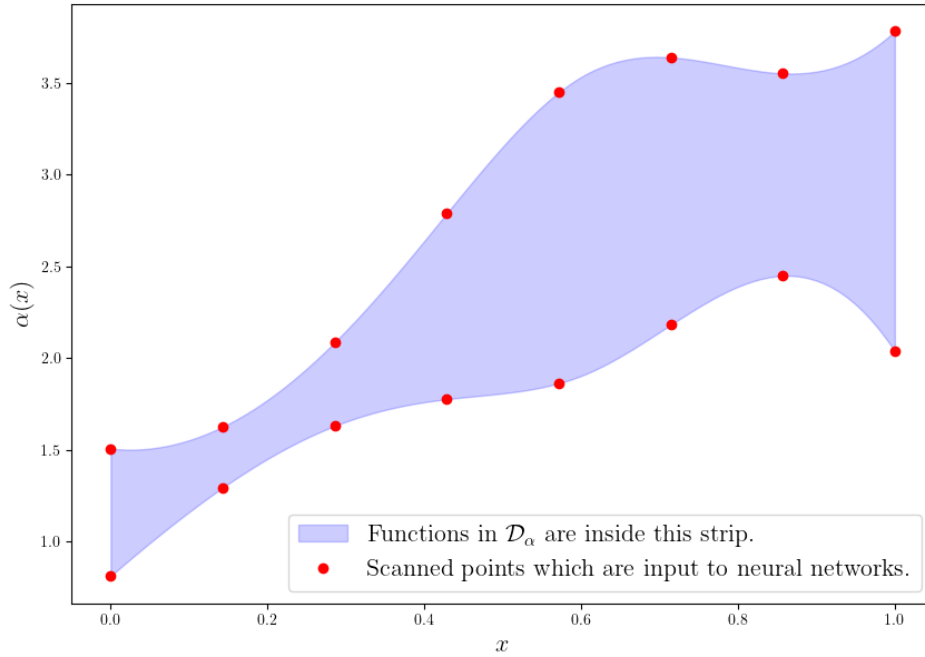


Figure 5. Example of what admissible sets in the test datasets can look like. In the test datasets, the admissible sets \mathcal{D}_α have a smooth edge and oscillate randomly as fast as chosen for the particular test dataset. The y -coordinates of the red dots are used as inputs for neural networks.

236 where \mathbf{x} is the vector of outputs of the neurons in the previous layer, \mathbf{w} are the weights of
 237 each connection from the previous layer to the given neuron, b is additional bias and the
 238 function a is the *activation function* of the neuron. The activation function may be one of
 239 many different functions but some common ones are the Sigmoid function $S(x) = \frac{1}{1+e^{-x}}$ and
 240 the ReLU (Rectified linear unit) function which is the piecewise defined function $R(x) = 0$
 241 when $x < 0$ and $R(x) = x$ when $x \geq 0$.

The machine learning part in using a neural network is when the weights w and biases b of each neuron are optimized on a given dataset. In such a task when the training dataset inputs are X and outputs are Y the problem is finding

$$\inf_{w,b} \|F(X, w, b) - Y\|,$$

242 where $F(X, w, b)$ is the output of the neural network for the entire training dataset with the
 243 weights and biases w, b . In practice the exact minimizer with respect to the weights and
 244 biases is not available. However, using numerical methods such as gradient descent such a
 245 minimization problem can be solved to a satisfactory level in many cases. For a more detailed
 246 introduction to the various different kinds of neural networks and the methods used to train
 247 them see for example [3].

248 **3.4. Architecture.** Four neural networks were created. One for each of the three training
 249 datasets and an extra one which uses all three training datasets combined as training data.

250 All networks have an input layer with 48 nodes. This corresponds to the admissible datasets of
 251 each parameter being scanned at eight points (see Figure 5). Implementation of the networks
 252 was done with Python version 3.8.12 [11] and Pytorch 1.9.0 [8]. Hyperparameter tuning was
 253 done with Optuna version 2.10.0 [1]. Optuna is a hyperparameter tuning package which
 254 uses various algorithms to scan the search space to find optimal parameters. In our case,
 255 the parameters being tuned are the number of hidden layers, number of nodes per hidden
 256 layer, activation function type, batch size and the hyperparameters of the Adam optimization
 257 algorithm. In Optuna, for sampling we picked the TPE (Tree-structured Parzen Estimator)
 258 algorithm, which is recommended in the documentation.

259 The choice of network architecture is a difficult open problem and our choices are based on
 260 intuition and trial and error. First, models with various structures were tested to give a broad
 261 idea of what kind of model could work and we found that relatively small networks with as few
 262 as two hidden layers and 20 neurons per layer are trained relatively well. We also tested ReLU
 263 and Sigmoid activation functions and found that both can give good results. In preliminary
 264 tests we used default settings of the Adam optimization algorithm but the hyperparameters
 265 of Adam will be tuned by Optuna. This gives us a good starting point to select a search space
 266 for Optuna. For the first three networks the search space given to Optuna was the following:

- 267 1. Number of hidden layers ≤ 3
- 268 2. Number of nodes in every hidden layer ≤ 100 (each hidden layer has the same amount
 269 of nodes)
- 270 3. Activation function of all nodes either Sigmoid or ReLU (all nodes have the same
 271 activation)
- 272 4. Batch size between 20 and 150.
- 273 5. Adam hyperparameters α (learning rate), β_1 , β_2 and ε :
 - 274 • $1 \times 10^{-5} \leq \alpha \leq 1 \times 10^{-2}$
 - 275 • $0.01 \leq \beta_1 \leq 0.9999$
 - 276 • $0.01 \leq \beta_2 \leq 0.9999$
 - 277 • $0 \leq \varepsilon \leq 1 \times 10^{-3}$

278 For the final network where we use all three datasets combined to train it the search space
 279 was adjusted to allow more hidden layers and nodes in each layer. The motivation for this is
 280 the fact that a larger dataset with more varied data should require a larger model to get a
 281 good training fit. The change done was the following:

- 282 1. Number of hidden layers ≤ 5
- 283 2. Number of nodes in every hidden layer ≤ 300 (each hidden layer has the same amount
 284 of nodes)

285 Each dataset was used over 2000 trials to create a neural network that works well for that
 286 specific dataset. The architectures that were found best by Optuna trials can be found in
 287 table 1. The first network for the simplest dataset (Network 1) is also a lot smaller in terms
 288 of neurons and layers than the latter ones with more complex training datasets. We can see
 289 that the level of complexity of the network increases when the training datasets complexity
 290 increases. For the last three networks the ReLU activation function outperforms Sigmoid
 291 however the differences are not massive so Sigmoid or some other nonlinear activation function
 292 may work reasonably well too.

Network	1	2	3	4
Activation	Sigmoid	ReLU	ReLU	ReLU
Hidden Layers	2	3	3	5
Nodes per layer	12	82	83	159
Batch Size	24	32	20	29
Learning Rate	0.0005	0.0068	0.0047	0.0052
β_1	0.7792	0.9018	0.7822	0.9151
β_2	0.8268	0.8606	0.4533	0.8770
ε	3.9×10^{-5}	2.9×10^{-4}	1.9×10^{-4}	3.9×10^{-4}

Table 1

Best parameter values for neural networks found for each dataset over 2000 Optuna trials. These values are used in the final trained neural networks. Network 1 was optimized while training with training dataset 1, Network 2 with training dataset 2, Network 3 with training dataset 3 and Network 4 with all three training datasets (1,2,3) combined.

293 **3.5. Tests.** Neural network performance is evaluated in two ways. The first one is to
 294 simply present the mean absolute error percentage (MAPE) which is defined as follows:

$$295 \quad \text{MAPE}(c, p) = \frac{100\%}{n} \sum_{t=1}^n \left| \frac{a_t - p_t}{a_t} \right|,$$

296

297 where c is the set of correct values, p is the set of predicted values and n is the amount of
 298 value-pairs.

299 The second way we evaluate performance is by comparing the networks outputs with
 300 known analytical upper and lower bounds. This comparison is done visually by plotting the
 301 correct approximation, predicted value and bounds for the test sets and the test portions of
 302 the training sets. These type of error bounds are well studied in [9] and [6]. The particular
 303 bounds that we use for problem (3.1-3.3) were derived in ([4], section 2.2).

304 For the analytical bounds we need some notations. Define

$$305 \quad c_1 := \min_{x \in (0,1)} \frac{\delta_1(x)}{\alpha_o(x) - \delta_1(x)}, \quad c_2 := \min_{x \in (0,1)} \frac{\delta_2(x)}{\beta_o(x) - \delta_2(x)},$$

306

307

$$308 \quad \overline{K} := \sqrt{1 + \max\{c_1, c_2\}} \quad \text{and} \quad \underline{K} := \sqrt{1 - \max\{c_1, c_2\}}.$$

310 The lower and upper bounds are defined in the next theorems.

311 **Theorem 3.1.** *The diameter of the solution set $\mathcal{S}(\mathcal{D})$ of the uncertain problem \mathcal{P} with ad-*
 312 *missible dataset \mathcal{D} has a guaranteed lower bound given by*

$$313 \quad \text{Diam}(\mathcal{S}(\mathcal{D})) \geq \underline{K} \frac{\int_0^1 \delta_1 |u'_o|^2 + \delta_2 |u_o|^2 + \delta_3 |u_o| \, dx}{\left(\|u_o\|_o^2 - \int_0^1 \delta_1 |u'_o|^2 + \delta_2 |u_o|^2 \, dx \right)^{\frac{1}{2}}},$$

314

Network	Test portion of training set	Test set 1	Test set 2	Test set 3
1	6.17(4.40)	7.94(6.71)	9.16(7.01)	37.08(7.33)
2	13.69(13.68)	10.97(6.31)	11.49(6.67)	30.86(10.63)
3	16.16(14.58)	11.07(7.01)	10.92(5.83)	38.86(14.92)
4	14.37(13.03)	7.58(5.32)	7.34(5.53)	45.09(11.36)

Table 2

Performance of each neural network on test sets expressed as the mean absolute error percentage (Standard Deviation) rounded to two decimal digits. Each network is evaluated on their respective training datasets test split which had 200 datapoints in the first column. In the three later columns are the performances of each network on the three testing datasets.

315 **Theorem 3.2.** *The diameter of the solution set $\mathcal{S}(\mathcal{D})$ of the uncertain problem \mathcal{P} with ad-*
316 *missible dataset \mathcal{D} has a guaranteed upper bound given by*

$$317 \quad \text{Diam}(\mathcal{S}(\mathcal{D})) \leq 2\bar{K} \left(\left(\int_0^1 \frac{(\delta_1 u'_o)^2}{\alpha_o - \delta_1} dx \right)^{\frac{1}{2}} + \bar{C}_F (|\delta_2| |u_o| + \delta_3) \right),$$
318

319 where $\bar{C}_F = \frac{1}{\pi} (\text{ess sup}_{\alpha \in \mathcal{D}_\alpha, x \in [0,1]} \alpha(x)^{-1})^{\frac{1}{2}}$.

320 The proofs of Theorems 3.1 and 3.2 with all the required preliminaries can be found in [4]
321 section 2. Creating such analytical bounds is based on a posteriori error estimates of the
322 functional type which are discussed at length in [9] and [6]. Since the proofs are rather long
323 and technical we omit them here.

324 **4. Results.** The performance of each neural network in terms of the MAPE value is
325 shown in table 2. In figures 6-8 the performance of the networks compared to the brute
326 force approximations and analytical bounds is compared using the test portion of the training
327 datasets. In figures 10-14 the networks performances are similarly showcased on the test
328 datasets.

329 In some cases the MAPE value is high but the figure still has some reasonability in what
330 the network outputs look like compared to the analytical bounds (e.g. Figure 14).

331 On the test portion of their respective training datasets network 1 has the highest accuracy.
332 This is expected since the dataset used to train network 1 is the most constrained. Looking at
333 the images (6)-(9) the performance of all four networks on the test sets related to their training
334 datasets seem decent. Each network outputs values which are always within the analytical
335 bounds. Network 4 is the only one with a few examples where the network output is rather
336 far away from the approximation but still within the bounds.

337 Performance on the three test sets varies. Network 1 and network 4 perform the best
338 on test set 1 in terms of the MAPE value. Network 4 keeps performing similarly well on
339 test set 2 while network 1 gets slightly worse. The added oscillation in test set 2 has this
340 effect since network 1 has not seen such highly oscillating data while network 4 has seen more
341 oscillating data (see Figures (10)-(12)). Networks 2 and 3 seem to suffer from the fact that
342 the range in which their training datasets are is larger than the test sets. This means there
343 are less training examples in the same range as the test set which explains why network 1

Dataset	Average time to create one datapoint
Training 1	1m 31s
Training 2	2m 28s
Training 3	6m 11s
Test 1	18m 27s
Test 2	29m 30s
Test 3	15m 42s

Table 3

Computational time to create datapoints in each training and testing dataset.

344 performs better. Still, both networks 2 and 3 are not completely lost on these test sets and
345 keep performing within the analytical bounds. From this we can tell that piecewise constant
346 oscillations approximate the smoother oscillations of the test sets reasonably well.

347 All four networks have a much higher error on test set 3 which shows us that the networks
348 do not easily generalize to examples where the admissible data has functions which are not
349 in the same range as the ones in the training examples. In Figure 14 the performance of
350 network 4 on test set 3 is shown. Here we can see that the network overestimates the value of
351 the diameter but there is still a linear growth which coincides with the growth of the correct
352 outputs. The network outputs are still within the analytical bounds.

353 Depending on the size of the neural network and luck the time to train them varied from
354 as little as five seconds to at most two minutes on a normal laptop. The total time to create
355 the optimised models using Optuna with 2000 trials took between five and ten hours for each
356 model. The computational time to create each dataset is shown in table 3. Training datasets
357 with their piecewise constant restrictions are generated more quickly in all cases. Test set 2
358 was much slower to create because of the higher oscillation allowed.

359 The amount of time saved when using a neural network rather than a Monte Carlo com-
360 putation is very large. Each of the neural networks created can be used to calculate an
361 approximation in less than one millisecond. The Monte Carlo method we implemented takes
362 an average of 56 seconds to create the solutions and another 188 seconds to compute the
363 maximum distance between the solutions. In total the time to use the Monte Carlo algorithm
364 to approximate the error quantity is 244 seconds.




Figure 6. Network 1 performance on test part of training dataset 1 which has 200 datapoints sorted by the ground truth approximation value. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.




Figure 7. Network 2 performance on test part of training dataset 2 which has 200 datapoints sorted by the ground truth approximation value. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.




Figure 8. Network 3 performance on test part of dataset 3 which has 200 datapoints sorted by the ground truth approximation value. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.

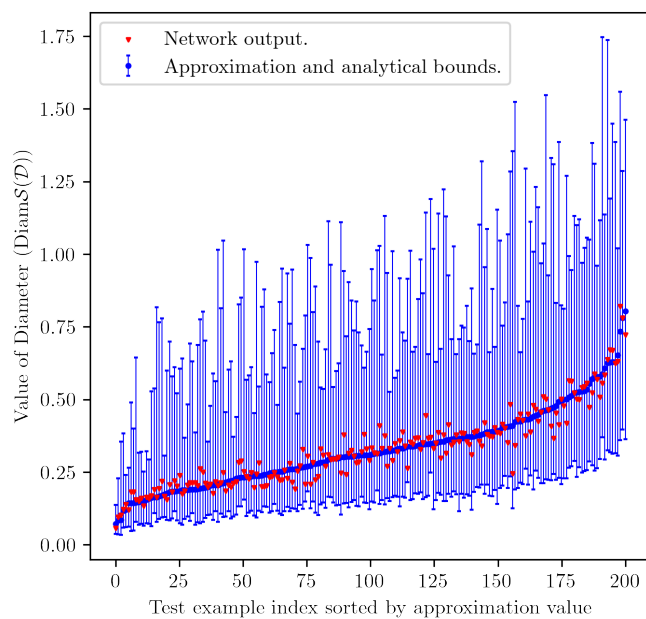


Figure 9. Network 4 performance on 200 random samples from the test part of the combined dataset sorted by the approximation values. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.

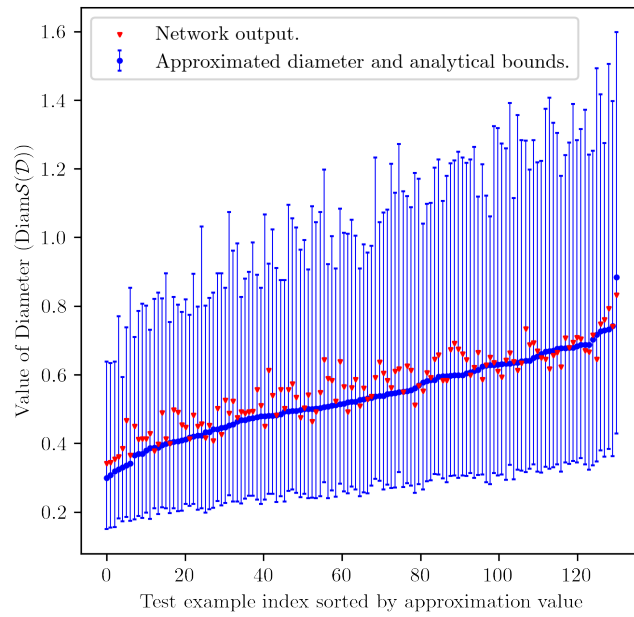


Figure 10. Network 1 performance on test set 1. Datapoints sorted by approximation values. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.

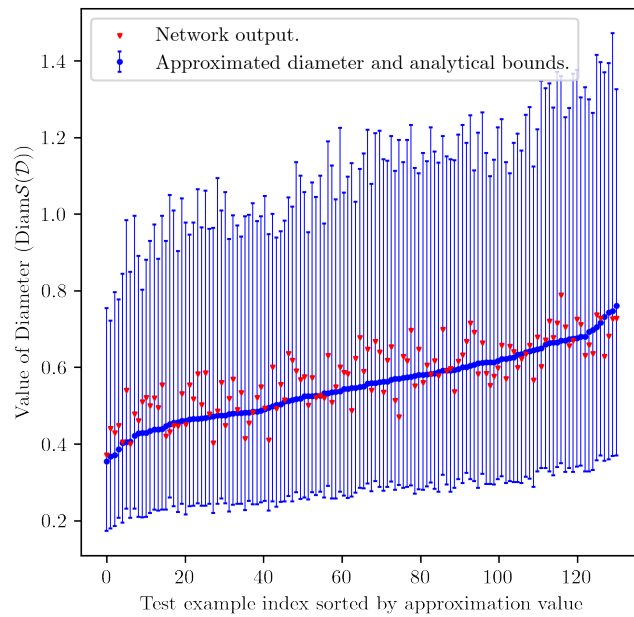


Figure 11. Network 1 performance on test set 2. Datapoints sorted by approximation values. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.




Figure 12. Network 4 performance on test set 1. Datapoints sorted by approximation values. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.




Figure 13. Network 4 performance on test set 2. Datapoints sorted by approximation values. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.




Figure 14. Network 4 performance on test set 3. Datapoints sorted by approximation values. Blue dots are the brute-force approximation of the diameter, error bars in blue are the analytical bounds given in 3.1 and 3.2, and red triangles are the neural network outputs.

365 **5. Discussion.** When trying to make an error quantifying machine learning model three
 366 things had to be considered. How to create data, model architecture and verifying model
 367 accuracy. In this paper the way we created data is the simplest method but it is also com-
 368 putationally expensive. Similar methods as the ones used to create data for training models
 369 which solve PDE's (see e.g. [7]) could be used to make this process quicker. In more diffi-
 370 cult problems this may even be practically necessary since such brute force methods may be
 371 unreasonably expensive.

372 When making data for a specific real-world problem one can more carefully select the type
 373 of training examples. The problem we discussed was not related to any physical values but
 374 only serves to show that such a model can perform in some constrained situation. It is most
 375 likely important that the training data coincides with the range of values that the practical
 376 application will deal with. Choosing very constrained datapoints for training data may help
 377 with quickly generating the dataset but accuracy may be sacrificed if this is done to too large
 378 an extent. In our example, using a rather constrained dataset still produced good results for
 379 some less constrained test datasets.

380 Dense neural networks worked well in our example but there is no reason to think that
 381 more complicated models like CNN's can't be used for quantifying uncertainty errors in some
 382 problem. For the example we provided even simpler regression models like nonlinear support
 383 vector machines could be tried. For more difficult problems it is hard to say how sophisticated
 384 the models required to tackle this problem are.

385 Accuracy verification of models may also be very difficult in higher dimensional nonlinear
 386 problems. For our example problem, both analytical and numerical methods can be used to

387 verify results but this is not always the case. In most real-world problems analytical methods
388 not at our disposal and one would like to avoid expensive computational methods as well.

389 Future research on this topic will be focused on creating machine learning models that es-
390 timate uncertainty errors of higher dimensional PDE's. A particularly interesting application
391 of uncertainty quantification models is in simulations. Scientific experiments in the future will
392 be done more and more by simulations and in order for a simulated experiment to be reliable
393 error quantification has to be addressed. Complex simulations by themselves are costly and
394 in such cases there are no efficient ways to quantify uncertainty errors quickly.

395 There are other interesting quantities that could be approximated with similar machine
396 learning models as the ones described. The error metric we used was the total error described
397 by the energy norm diameter of the solution set but depending on the use-case this is not
398 always the most valuable information. In certain applications for example local errors are more
399 useful than the global error. For such problems a model should be trained to approximate a
400 local error metric instead.

401 **6. Conclusions.** The experiments show that a neural network can be used to accurately
402 approximate errors caused by uncertainty in the case of a linear differential equation. In all of
403 the tests conducted the neural networks work well enough to be considered as an option instead
404 of previously known analytical methods. These results show that if a suitable training dataset
405 can be created it is feasible to use machine learning models for uncertainty quantification.
406 Compared to the brute-force approximation and known analytical bounds the models perform
407 with similar accuracy. The approximation is essentially a Monte Carlo type quantity which
408 means there is hope of replacing this type of slow and expensive methods with faster machine
409 learning model inference. In terms of time saving the neural networks are much faster to use
410 than a Monte Carlo method. For this simple linear problem the Monte Carlo approximation
411 still takes around four minutes to compute while the neural network is inferenced in less than
412 one millisecond.

413 In our example, the training datasets and test datasets had some differences. Even though
414 the test data was more general, the neural networks performed in a good way. It is very useful
415 if training examples can be picked from a more constrained class of examples since in some
416 cases this makes the generation of the training example much more computationally fast. The
417 only test set on which the models did not perform very well is the last one. In that case, the
418 range of values for functions in the admissible data of test examples was outside of the range
419 in which they were for training datasets. Higher oscillation of the admissible datasets in test
420 examples did not have a catastrophic effect.

421 Creation of networks and data took a fair amount of time even for this simple problem.
422 The data was generated on a normal laptop and it took around two minutes to generate a single
423 training datapoint. The datapoints in the test datasets took longer at around ten minutes
424 for a single datapoint. It is clear that the process of data generation for higher dimensional
425 problems will have to be done by organisations with large computational resources or by using
426 more clever methods which require less computation.

427 Doing the same thing for 2-D and 3-D problems will be computationally more difficult
428 for both creation of data and training models. However neural networks even when they are
429 much more complicated than the ones we have used in this paper are trained with relatively

430 small amounts of computational resources. The main computational time problem will be
431 with creating the data. In this paper the Monte Carlo method used required creating 2000
432 different solutions for a simple ODE and that took one minute or so. For a complicated
433 non-linear high dimensional PDE generating this set of solutions may take a very long time
434 depending on what type of solvers can be used for that problem. Computing the maximum
435 distance between solutions will also be more expensive when more dimensions are introduced
436 since numerical integration gets more expensive with more dimensions. It may also require
437 more datapoints to have a suitable training dataset for a network to learn properly. It is
438 reasonable to expect that the more complicated the PDE model and the more general the
439 restrictions on the dataset are the more training data is required.

440 **Acknowledgments.** Thanks to Professor Sergei Repin for fruitful discussions and contri-
441 butions.

442

REFERENCES

- 443 [1] T. AKIBA, S. SANO, T. YANASE, T. OHTA, AND M. KOYAMA, *Optuna: A next-generation hyperparameter*
444 *optimization framework*, in Proceedings of the 25rd ACM SIGKDD International Conference on
445 Knowledge Discovery and Data Mining, 2019.
- 446 [2] P. G. CIARLET, *The finite element method for elliptic problems*, SIAM, 2002.
- 447 [3] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep Learning*, MIT Press, 2016. [http://www.](http://www.deeplearningbook.org)
448 [deeplearningbook.org](http://www.deeplearningbook.org).
- 449 [4] V. HALONEN, *Accuracy analysis of uncertain variational problems with analytical and machine learning*
450 *methods*, University of Jyväskylä, 2021. Master's Thesis.
- 451 [5] I. HLAVACEK, J. CHLEBOUN, AND I. BABUSKA, *Uncertain input data problems and the worst scenario*
452 *method*, Elsevier, 2004.
- 453 [6] O. MALI, P. NEITTAANMÄKI, AND S. REPIN, *Accuracy verification methods: Theory and algorithms*,
454 vol. 32, Springer Science & Business Media, 2013.
- 455 [7] A. MUZALEVSKIY, P. NEITTAANMÄKI, AND S. REPIN, *Generation of Error Indicators for Partial Dif-*
456 *ferential Equations by Machine Learning Methods*, Springer International Publishing, Cham, 2022,
457 pp. 63–96, https://doi.org/10.1007/978-3-030-70787-3_6.
- 458 [8] A. PASZKE, S. GROSS, F. MASSA, A. LERER, J. BRADBURY, G. CHANAN, T. KILLEEN, Z. LIN,
459 N. GIMELSHEIN, L. ANTIGA, A. DESMAISON, A. KOPF, E. YANG, Z. DEVITO, M. RAISON,
460 A. TEJANI, S. CHILAMKURTHY, B. STEINER, L. FANG, J. BAI, AND S. CHINTALA, *Pytorch: An*
461 *imperative style, high-performance deep learning library*, in Advances in Neural Information Pro-
462 cessing Systems 32, Curran Associates, Inc., 2019, pp. 8024–8035, [http://papers.neurips.cc/paper/](http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf)
463 [9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf](http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf).
- 464 [9] S. I. REPIN, *A posteriori estimates for partial differential equations*, in A Posteriori Estimates for Partial
465 Differential Equations, de Gruyter, 2008.
- 466 [10] G. SCHUËLLER, *A state-of-the-art report on computational stochastic mechanics*, Probabilistic Engineering
467 Mechanics, 12 (1997), pp. 197–321, [https://doi.org/https://doi.org/10.1016/S0266-8920\(97\)00003-9](https://doi.org/https://doi.org/10.1016/S0266-8920(97)00003-9),
468 <https://www.sciencedirect.com/science/article/pii/S0266892097000039>.
- 469 [11] G. VAN ROSSUM AND F. L. DRAKE, *Python 3 Reference Manual*, CreateSpace, Scotts Valley, CA, 2009.