

DeepONet-grid-UQ: A trustworthy deep operator framework for predicting the power grid's post-fault trajectories

C. Moya, M. Yue

To be published in "NEUROCOMPUTING"

May 2023

Interdisciplinary Science Department
Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Energy Efficiency and Renewable Energy (EERE), Renewable Power Office

Notice: This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DeepONet-Grid-UQ: A Trustworthy Deep Operator Framework for Predicting the Power Grid’s Post-Fault Trajectories

Christian Moya, Shiqi Zhang, Meng Yue, and Guang Lin

February 16, 2022

Abstract

This paper proposes a new data-driven method for the reliable prediction of power system post-fault trajectories. The proposed method is based on the fundamentally new concept of Deep Operator Networks (DeepONets). Compared to traditional neural networks that learn to approximate functions, DeepONets are designed to approximate nonlinear operators. Under this operator framework, we design a DeepONet to (1) take as inputs the fault-on trajectories collected, for example, via simulation or phasor measurement units, and (2) provide as outputs the predicted post-fault trajectories. In addition, we endow our method with a much-needed ability to balance efficiency with reliable/trustworthy predictions via uncertainty quantification. To this end, we propose and compare two methods that enable quantifying the predictive uncertainty. First, we propose a *Bayesian DeepONet* (B-DeepONet) that uses stochastic gradient Hamiltonian Monte-Carlo to sample from the posterior distribution of the DeepONet parameters. Then, we propose a *Probabilistic DeepONet* (Prob-DeepONet) that uses a probabilistic training strategy to equip DeepONets with a form of automated uncertainty quantification, at virtually no extra computational cost. Finally, we validate the predictive power and uncertainty quantification capability of the proposed B-DeepONet and Prob-DeepONet using the IEEE 16-machine 68-bus system.

1 Introduction

Many critical infrastructure systems depend heavily on a robust, reliable, and always-on power grid. However, our power grid is constantly exposed to rare but severe events and disturbances that could compromise its operation and security. In the worst-case scenario, such disturbances could lead to instabilities, which, in turn, could trigger system-wide blackouts.

To assess the power grid’s dynamic security [17], system operators implement an offline transient analysis procedure known in practice as the dynamic security analysis (DSA), usually based on the $N-1$ criteria [1]. The $N-1$ based DSA (1) seeks to predict whether the power grid will remain safely operating after facing a significant disturbance (*e.g.*, a transmission line fault or a generator trip) selected from a set of credible disturbances and (2) requires simulating the power grid’s dynamic response.

To simulate the power grid’s dynamic response, one must simulate a set of nonlinear differential-algebraic equations (DAE) [17]. Simulating this set of DAEs is, however, a very challenging task. Indeed, the classical explicit integration methods fail catastrophically on such a task [13]. As a result, most of the software tools used for transient stability implement numerically stable integration methods [17] designed specifically for DAEs. However, the computational cost and memory required for these methods are very high and constitute the main obstacle towards deploying a more efficient and possibly online transient analysis [25]. With the transformation

the power grid now faces (*e.g.*, by deploying a large number of distributed renewable energy resources or a liberalized market), soon, it will become imperative for electric utilities to assess its dynamic security near real-time, which calls for more efficient transient analysis methods.

1.1 Prior works

Of course, to accelerate the simulation of DAEs, we can resort to schemes aiming at parallelizing the existing numerical methods. Parallel computing, however, consumes a high amount of resources, and the performance can quickly reach the saturation, which could prevent its application in near-real-time transient analysis. Besides parallel computing, other methods have been proposed to assess transient stability efficiently in the last decades. Here, we classify them into (1) transient analysis methods that exploit the mathematical structure of the power grid models and (2) Machine Learning (ML)-enabled transient analysis methods.

1.1.1 Power grid model-based transient analysis methods

An alternate approach to the traditional offline transient analysis method was proposed in the early '80s under the name of *direct methods* [23, 26, 12] (or energy function-based methods). These methods provide transient stability certification without the time-consuming numerical simulation of the post-fault dynamics. More specifically, direct methods assess transient stability by analyzing the power grid dynamical model with the aid of an energy function (*i.e.*, a Lyapunov-like function), which certifies convergence of the states to the stable equilibrium point. Modern versions of direct methods [7, 6] have been successfully engineered to the point that they were implemented at the utility level. However, the scalability and conservativeness of the classical energy-function method limit their applicability to even relatively large power grids.

Other works have also addressed the transient analysis problem using power grid dynamical models. For example, Dorfler *et al.* [9] studied the stability of power grids with overdamped generators via studying the synchronization of Kuramoto oscillators. The authors of [15] used Hamilton Jacobi reachability to find an inner approximation of the region of attraction, which enables the online transient stability certification of initial conditions. Finally, [2] proposed using the sum of squares programming, and [3] proposed using a port Hamiltonian approach. All the above methods use simplified power grid dynamics. So, their applicability to more complex power grids remains unclear. In addition, most of these works produce only stability margins but not the power grid post-fault trajectories, which system operators and planners often need.

1.1.2 Machine Learning (ML)-enabled transient analysis methods

Machine Learning holds the promise to revolutionize science and engineering, including power engineering. Thus, as expected, many works have proposed to address the transient analysis problem using Machine Learning and Artificial Intelligence-inspired methodologies. The main idea behind most of these works is to learn a mapping from initial states/responses, collected via simulation or phasor measurement units (PMU), to a certificate of transient stability.

For example, He *et al.* [11] proposed using ensemble tree learning to construct a hierarchical characterization of the dangerous regions in the space of possible contingencies and operating/initial states. In [28], the authors proposed a classification approach to learn the decision rules for certifying transient stability from initial responses. Using recurrent neural networks, the authors in [14] proposed a self-adaptive transient analysis method. In [10], the authors developed

a convolutional neural network that takes as the input measurements collected using PMUs and outputs a transient stability certificate. However, the main drawback of these aforementioned ML-based approaches is that they focus primarily on developing a binary/discrete indicator that certifies whether the power grid is transient stable or not after a disturbance.

To provide a more quantitative metric for transient stability, Zhu *et al.* [31] developed a convolutional neural network that takes as the input the fault-on trajectory collected using PMUs and outputs not only a transient stability certificate but also an estimate of the stability margin. This ML-based tool can be fast, but it may be insufficient for system operators and planners. These operators and planners may be interested in knowing the trajectories of various states of the power grid after the disturbance. These trajectories are crucial to estimate whether any voltage or frequency will violate pre-defined limits and trigger, for example, load shedding. Hence, in the authors' previous work [18], we developed a deep-learning framework that uses the long-short-term memory (LSTM) to predict the transient responses of power grid states. The effectiveness and speed of this approach in [18] was undoubtedly a first step towards the development of a more complete and efficient ML-based transient analysis tool for power grids. However, our previous work, like most of the proposed ML-based transient analysis works, (1) requires a considerable amount of supervision for successful deployment and (2) fails to tackle the inherent trade-off between being efficient and being reliable.

1.2 Our work

This paper proposes a machine learning-based method for predicting the power grid's post-fault transient response to faults. However, compared to existing ML-based methods, we employ a different class of neural networks, introduced in the seminal work of Lu *et al.* [19], that we train to approximate the operator arising from such transient response. Under this operator framework, our objectives are:

1. *Post-fault trajectory prediction*: deriving a Deep Operator Network (DeepONet) framework that observes the fault-on trajectory (of a given state) to predict the post-fault trajectory (of the same state).
2. *Trustworthy prediction*: designing an effective strategy that attaches a measure of confidence/reliability to the predicted trajectory.

Compared to [18], we do not focus on improved mean trajectory prediction. One can always improve the mean predictive performance by using novel neural network architectures or enriching the datasets (which may not be possible in practice). Hence, in this paper, we consider a smaller dataset than in [18]. Furthermore, during training, we will not use information from neighboring buses, which generally improves the mean prediction. The above two assumptions make the proposed problem of providing trustworthy/reliable post-fault trajectory predictions much more challenging. Let us conclude this section by enumerating our contributions and describing the paper's organization.

1. We first design (in Section 2) a data-driven deep operator network (DeepONet) framework for post-fault trajectory prediction. Our DeepONet learns to approximate the solution operator that maps trajectories containing fault-on information to their corresponding post-fault trajectories. Once trained, the proposed DeepONet provides a more efficient way of calculating the post-fault trajectories than traditional transient response simulators.
2. We then endow the proposed DeepONet with the ability to provide trustworthy predictions.

To this end, we develop two methods that quantify the heteroscedastic aleatoric uncertainty (*i.e.*, the input and data-inherent uncertainty) of the predicted trajectories.

3. In the first method, we design (in Section 3) a Bayesian framework for DeepONets. Here the Bayesian DeepONet (B-DeepONet) represents the prior of the trainable parameters. We then employ the (theoretically sound) Stochastic Gradient Hamiltonian Monte-Carlo (SGHMC) algorithm to sample the corresponding posterior. We use the collected samples to quantify the uncertainty of predicted post-fault trajectories.
4. For our second method, we propose (in Section 4) a probabilistic DeepONet (Prob-DeepONet) framework. This framework uses a special DeepONet architecture and probabilistic training to transform the DeepONet predictions into point-wise Gaussian estimates. Predicting Gaussian estimates enables quantifying the uncertainty of the predicted trajectories at virtually no extra computational cost.
5. Finally, we test (in Section 5) the trustworthiness and predictive power of the proposed B-DeepONet and Prob-DeepONet using the IEEE-68-bus power grid.

We organize the rest of this paper as follows. In Section 2, we formulate the problem of predicting the post-fault trajectories as an operator regression problem. Moreover, we introduce the DeepONet to approximate such an operator. The B-DeepONet and Prob-DeepONet frameworks that enable trustworthy post-fault trajectory prediction are detailed, respectively, in Section 3 and Section 4. In Section 5, we test, using the IEEE-69 bus power grid, the reliability and predictive power of the proposed frameworks. Section 6 discusses our results and future work. Finally, Section 7 concludes the paper.

2 Operator Regression for Transient Response Prediction

This section describes how to pose the post-fault trajectory prediction problem as an operator regression problem.

Post-fault trajectory prediction: Faults and disturbances on transmission lines (or other large grid equipment) can drive the power grid to instability. In a typical fault scenario, the power grid experiences the fault and moves away from the pre-fault stable equilibrium until the system clears the fault [6]. Once the fault is cleared, the power grid experiences a transient response known as the post-fault dynamics. For power grid operators, it is then essential to know whether these post-fault dynamics converge or not to a stable equilibrium point [17]. If the trajectories fail to converge, a series of conventional protection schemes will be triggered, leading, for example, to load shedding.

Operator regression: We formalize next how to predict post-fault trajectories using operator regression. Assume the power grid was operating at a stable equilibrium point. Then, at some time t_f , the grid experiences a fault, which results in a structural change due to the system protective actions [6]. Assume the resulting fault-on dynamics are confined to the time interval $[t_f, t_{cl}]$, where t_{cl} denotes the time when the fault is cleared. During $[t_f, t_{cl}]$, we can describe the fault-on dynamics using the set of differential-algebraic equations (DAE):

$$\dot{x}(t) = f_F(x(t), z(t)) \quad t \in [t_f, t_{cl}], \quad (1a)$$

$$0 = g_F(x(t), z(t)), \quad (1b)$$

where x is the vector of dynamic states and z the vector of algebraic states. Now, assume the system clears the fault at time t_{cl} and no additional protective actions occur after t_{cl} . Then, the

post-fault dynamics, *i.e.*, the dynamics within the interval $(t_{cl}, T]^1$, are given by the DAEs:

$$\dot{x}(t) = f_{PF}(x(t), z(t)) \quad t \in (t_{cl}, T], \quad (2a)$$

$$0 = g_{PF}(x(t), z(t)). \quad (2b)$$

In this paper, we aim to learn the operator, denoted as $G^\dagger$, mapping a fault-on trajectory (of a selected dynamic or algebraic state $w \in \{x, z\}$), which satisfies (1), to the post-fault trajectory (of the *same* state w), which satisfies (2), *i.e.*,

$$G^\dagger : w|_{[t_f, t_{cl}]} \mapsto w|_{(t_{cl}, T]}. \quad (3)$$

Here $w|_{[t_f, t_{cl}]}$ (resp. $w|_{(t_{cl}, T]}$) denotes the fault-on (resp. post-fault) trajectory of the selected dynamic/algebraic state w .

Remark. On observing the fault-on trajectory. Observe that in (3) the operator $G^\dagger$ takes as the input the exact fault-on trajectory, *i.e.*, $w|_{[t_f, t_{cl}]}$. While this may be possible when using this operator setting for offline post-fault trajectory prediction (*i.e.*, using numerical simulation), in the online setting, it may not be possible to exactly capture $w|_{[t_f, t_{cl}]}$ using PMUs. As a result, we propose next a relaxed operator regression problem for post-fault trajectory prediction.

We start by splitting the time domain of transient stability, denoted as $[t_0, T]$, into two adjacent regions. The first region, denoted as T_u , satisfies $[t_f, t_{cl}] \subseteq T_u$ and corresponds to the time interval containing the domain of all possible fault-on trajectories. The second region, denoted as T_G , satisfies $T_G = [0, T] \setminus T_u$ and corresponds to the time domain of post-fault trajectories. Thus, we define the “relaxed” operator $G^\dagger$ for post-fault trajectory prediction as:

$$G^\dagger : w|_{T_u} \mapsto w|_{T_G}. \quad (4)$$

To learn the operator (4), we use the Deep Operator Network (DeepONet) framework introduced in [19] based on the universal approximation theorem to nonlinear operators [4].

2.1 Review of DeepONet

We now review the original formulation of DeepONet [19]. Let $G^\dagger$ denote a nonlinear operator. This operator $G^\dagger$ maps an input function u to an output function $G^\dagger(u)$. Let $y \in Y$ denote a point in the output function domain Y (often a subset of $\mathbb{R}^d$). Then, the goal of the DeepONet G_θ , with trainable parameters $\theta \in \mathbb{R}^p$, is to approximate the operator $G^\dagger(u)(y)$ at $y \in Y$. To this end, the DeepONet G_θ uses the neural network architecture depicted in Figure 1, consisting of two sub-networks referred to as the Branch Net and the Trunk Net.

The *Branch* Net processes the input function information. Let $(x_1, \dots, x_m)$ denote points in the domain of u (we will refer to these points as the *sensors*), such that $(u(x_1), \dots, u(x_m))$ is a discrete representation of the input u . The Branch Net takes this discretized u as the input and outputs a vector of features $b \in \mathbb{R}^q$. On the other hand, the *Trunk* Net processes points in the domain Y of the output function. To this end, the Trunk Net takes $y \in Y$ as the input and outputs a vector of features $\tau \in \mathbb{R}^q$. Note that since the Trunk Net’s output τ solely depends on the input coordinates y , it is natural to interpret the components of τ as a collection of basis functions defined on Y , *i.e.*,

$$\tau = (\varphi_1(y), \dots, \varphi_m(y)).$$

¹ T denotes a time (in the order of seconds) that can capture the power grid’s transient response.




Figure 1: The vanilla DeepONet architecture. The crossed node on the right indicates the DeepONet output, which we obtain by taking the inner product between the output features of the Branch (b) and Trunk (τ) Nets.

The output of the DeepONet then combines the output features from the Branch Net b and the Trunk Net τ using an inner product:

$$G_{\theta}(u(x_1), \dots, u(x_m))(y) := \langle b, \tau \rangle + \tau_o = \sum_{i=1}^q b_i \cdot \varphi_i(y) + \tau_o, \quad (5)$$

where τ_o is a trainable parameter added to the DeepONet's output. From the above, one can interpret the output of the Branch Net b as the trainable coefficients for the trainable basis functions τ produced by the Trunk Net. To simplify our notation, in the rest of this work, we omit writing the DeepONet explicit dependency on the discretized input and use the simplified notation $G_{\theta}(u)(y)$.

Finally, we train the DeepONet G_{θ} to approximate the nonlinear solution operator $G^{\dagger}$ by minimizing a mean square loss function on the training dataset $\mathcal{D} = \{u_i, y_i, G^{\dagger}(u_i)(y_i)\}_{i=1}^N$, *i.e.*,

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \left| G_{\theta}(u_i)(y_i) - G^{\dagger}(u_i)(y_i) \right|^2. \quad (6)$$

The DeepONet for post-fault trajectory prediction: Under the notation introduced in Section 2.1, the DeepONet that approximates the operator (4) for post-fault trajectory prediction (see Figure 2) is given as follows. The sensors $(x_1, \dots, x_m)$ correspond to sampling times $(t_1, \dots, t_m)$ within the time domain T_u , *i.e.*, $t_i \in T_u$ for all $i = 1, \dots, m$. The input function corresponds to the trajectory containing fault-on information, *i.e.*, $u \equiv w|_{T_u}$. Moreover, the output function corresponds to the post-fault trajectory, *i.e.*, $G^{\dagger}(u) \equiv w|_{T_G}$. Thus, the domain of the output function corresponds to the time domain of post-fault trajectories, *i.e.*, $Y \equiv T_G$, and y denotes a time sampled from T_G . Throughout the rest of this paper, we keep the notation introduced in Section 2.1 that follows [19]. This is because the methods that we develop next for trustworthy post-fault trajectory prediction can also be used for other operator learning problems.

We conclude this section with the following remark. DeepONets have shown remarkable approximation and generalization capabilities. Furthermore, DeepONets can learn operators efficiently with relatively small datasets. However, DeepONets inherits some of the problems from more traditional neural networks. Their reliability deteriorates when the input training data are noisy or do not cover all possible scenarios. One should note that noisy and incomplete training data is inherent in all power grid prediction problems. To overcome this reliability issue, we develop




Figure 2: The DeepONet for post-fault trajectory prediction. The sensors are $(t_1, \dots, t_m)$, where $t_i \in T_u$. The input function is $u = w|_{T_u}$. The output function is $G^\dagger(u) = w|_{T_G}$. Thus the goal of DeepONet G_θ is to approximate $G^\dagger(u)(y)$ at a point $y \in T_G$.

in the following sections two methods that enable DeepONets to provide trustworthy predictions through uncertainty quantification.

3 Bayesian Deep Operator Network

In this section, we propose a *Bayesian DeepONet* (B-DeepONet) framework (see Figure 3) to predict post-fault trajectories and quantify the aleatoric uncertainty. The proposed B-DeepONet consists of three components. (1) We represent the *prior* $p(\theta)$ of the trainable parameters $\theta \in \mathbb{R}^p$ using the Bayesian DeepONet; (2) We compute the *likelihood* $p(\mathcal{D}|\theta)$ using the forward pass of DeepONets and the measurements $G^\dagger(u)(y)$; and (3) we sample from the *posterior* $p(\theta|\mathcal{D})$ using the (theoretically sound) Stochastic Gradient Hamiltonian Monte-Carlo algorithm. We review next the Hamiltonian Monte-Carlo (HMC) algorithm [22].

3.1 Hamiltonian Monte-Carlo (HMC)

The goal of HMC [22] is to sample from the posterior distribution of the trainable parameters θ given a set of independent observations $d \in \mathcal{D}$. This target posterior is:

$$p(\theta|\mathcal{D}) \propto \exp(-U(\theta)),$$

where U is the potential energy defined as

$$U(\theta) = - \underbrace{\sum_{d \in \mathcal{D}} \log p(d|\theta)}_{\text{likelihood}} - \underbrace{\log p(\theta)}_{\text{prior}}. \quad (7)$$

To generate the samples, HMC uses the Hamiltonian system:

$$H(\theta, r) = U(\theta) + \frac{1}{2} r^\top M^{-1} r.$$

Here, $r \in \mathbb{R}^p$ is a vector of auxiliary momentum variables, and M (often set to be the identity matrix) is the mass matrix. The samples generated using $H(\theta, r)$ follow the joint distribution of (θ, r) given by

$$\pi(\theta, r) \propto \exp(-H(\theta, r)).$$




Figure 3: The Bayesian DeepONet framework. The Bayesian DeepONet represents the prior for the trainable parameters $\theta \in \mathbb{R}^p$. One computes the likelihood using the DeepONet forward pass and the training targets. Finally, one estimates the posterior using stochastic gradient Hamiltonian Monte-Carlo (HMC).

If we discard the obtained r samples, then the θ samples have as marginal distribution the desired target distribution $p(\theta|\mathcal{D})$ [22]. To propose new samples, HMC simulates the following Hamiltonian dynamics:

$$d\theta = M^{-1}r \, dt \quad (8a)$$

$$dr = -\nabla U(\theta) \, dt. \quad (8b)$$

To simulate from (8), the HMC algorithm uses, for *e.g.*, the *leapfrog* discretization method. In addition, a *Metropolis-Hasting* step is used to accept/reject samples, which can contain inaccuracies as a result of the leapfrog discretization.

The main disadvantage of HMC is the computational cost we require to (1) compute the gradient $\nabla U(\theta)$ when the dataset $\mathcal{D}$ is large and (2) perform the Metropolis-Hasting step. To tackle this computational cost drawback, we propose using *Stochastic Gradient* HMC [5] for sampling the posterior distribution of the DeepONet parameters θ .

3.2 Stochastic Gradient HMC (SGHMC)

In SGHMC [5], instead of computing the expensive gradient $\nabla U(\theta)$, which requires traversing the entire dataset $\mathcal{D}$, SGHMC considers a noisy estimate of this gradient based on a mini-batch $\tilde{\mathcal{D}} \subset \mathcal{D}$ sampled uniformly at random, *i.e.*,

$$\nabla \tilde{U}(\theta) = -\frac{|\mathcal{D}|}{|\tilde{\mathcal{D}}|} \sum_{d \in \tilde{\mathcal{D}}} \nabla \log(d|\theta) - \nabla \log p(\theta). \quad (9)$$

If the d 's are independent, then via the central limit theorem, the noisy gradient is given approximately by [5]:

$$\nabla \tilde{U}(\theta) \approx \nabla U(\theta) + \mathcal{N}(0, V(\theta)),$$

Algorithm 1: Stochastic Gradient HMC

```
1 Require: initial state  $\theta^0$  and time step size  $\epsilon_t$ ; for  $k = 1, \dots, N$  do
2   sample  $r^{k-1} \sim \mathcal{N}(0, M)$ 
3    $(\theta_0, r_0) = (\theta^{k-1}, r^{k-1})$ 
4   for  $i = 1, \dots, m$  do
5     simulate equation (10):
6      $\theta_i = \theta_{i-1} + \epsilon_t M^{-1} r_{i-1}$ 
7      $r_i = r_{i-1} - \epsilon_t \nabla \tilde{U}(\theta_i) - \epsilon_t C M^{-1} r_{i-1} + \mathcal{N}(0, 2(C - \hat{B})\epsilon_t)$ 
8   end
9    $(\theta^k, r^k) = (\theta_m, r_m)$ 
10 end
11 Calculate  $\{\{G_{\theta^{N+1-k}}(u)(y) : y \in Y_{\text{mesh}}\}_{k=1}^M$  as predicted samples of the true post-fault trajectory  $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$ 
```

where $V(\theta)$ is the covariance of the stochastic gradient noise,² which could depend on θ and $|\tilde{\mathcal{D}}|$.

To simulate samples using SGHMC, we use the following modified version of the dynamics (8) that adds a *friction* term to the momentum update, *i.e.*,

$$d\theta = M^{-1} r \, dt \quad (10a)$$

$$dr = -\nabla \tilde{U}(\theta) \, dt - C M^{-1} r \, dt + \mathcal{N}(0, 2(C - \hat{B})\epsilon_t) + \mathcal{N}(0, 2B\epsilon_t). \quad (10b)$$

Here $C \succeq \hat{B}$ is a user-specified *friction* term and $\hat{B}$ is an estimate of the noise model $B(\theta) = \frac{1}{2\epsilon} V(\theta)$, where ϵ is the step size of the discretization scheme. As demonstrated in [5], the above dynamics, with the proper selection of C and $\hat{B}$, yields the desired stationary distribution $\pi(\theta, r) \propto e^{-H(\theta, r)}$. We provide the details for implementing SGHMC in Algorithm 1.

To provide a trustworthy prediction of the post-fault trajectories (see Algorithm 1 for details), denoted as $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$, for a given input u and over the given mesh of points $y \in Y_{\text{mesh}} \subset Y \equiv T_G$, we must sample from the posterior $p(\theta|\mathcal{D})$ and obtain the M -ensemble of DeepONets with parameters $\{\theta^{N+1-k}\}_{k=1}^M$ depicted in Figure 4. Then, using this ensemble, we can simulate M trajectories $\{G_{\theta^{N+1-k}}(u)(y) : y \in Y_{\text{mesh}}\}_{k=1}^M$ and compute their statistics. Here, we compute the mean and standard deviation. We use the mean to predict the true solution $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$ and the standard deviation to quantify the uncertainty.

One should note that trustworthy post-fault trajectory prediction requires simulating M trajectories $\{G_{\theta^{N+1-k}}(u)(y) : y \in Y_{\text{mesh}}\}_{k=1}^M$ using the M -ensemble of DeepONets produced by Algorithm 1. However, simulating these many trajectories for large-scale power grids might become computationally expensive, which precludes the application of the proposed B-DeepONet to online post-fault trajectory prediction. While some of the computational costs may be alleviated using a parallel computing framework, we believe B-DeepONets are more suitable for offline post-fault trajectory prediction.

4 Probabilistic Deep Operator Network

The main drawback of most data-driven methods for transient stability is that they fail to handle the inherent trade-off between providing (1) trustworthy predictions and (2) efficient/fast

²We have adopted the same notation as in [5], that is $\mathcal{N}(\mu, \Sigma)$ represents a random variable distributed according to the multivariate Gaussian.




Figure 4: M -ensemble of DeepONet parameters $\{\theta^{N+1-k}\}$ sampled from the posterior $p(\theta|\mathcal{D})$ using SGHMC. We use this ensemble to simulate post-fault trajectories $\{G_{\theta^{N+1-k}}(u)(y) : y \in Y_{\text{mesh}}\}_{k=1}^M$. The mean of these trajectories is used to predict the true trajectory $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$. The standard deviation is used to quantify the aleatoric uncertainty.

predictions. Our proposed B-DeepONet (see Section 3) gave us the tools (supported by the solid mathematical results detailed in [5]) to provide reliable predictions by estimating their aleatoric uncertainty. However, the increased predictive reliability gained by B-DeepONets comes at the price of decreased predictive efficiency.

Motivated by the need for both trustworthy and efficient post-fault trajectory predictions, here, in this section, we take a more *empirical* approach and design a probabilistic training framework for DeepONets, which we refer to as *Prob-DeepONet*. Prob-DeepONet was introduced in one of the authors' previous works [29] and will enable us to provide efficient point-wise aleatoric uncertainty estimates of post-fault trajectory predictions.

Prob-DeepONet requires building a deep operator network architecture (depicted in Figure 5) with the ability to learn estimates of aleatoric uncertainty. As described in [29], this can be achieved by casting the DeepONet outputs as Gaussian posterior distributions. We select the outputs to represent Gaussian distributions because we have observed (see Appendix 1 for more details) that the network errors $(G_\theta(u)(y) - G^\dagger(u)(y))$ experienced by DeepONets when trained via gradient descent algorithms (*e.g.*, Adam) follow a Gaussian distribution. More specifically, we build a DeepONet architecture that predicts the point-wise statistics $\mu_\theta(u)(y)$ and $\sigma_\theta(u)(y)$ of the Gaussian distribution $\mathcal{N}(\mu_\theta(u)(y), \sigma_\theta^2(u)(y))$ of the *true* target operator values $G^\dagger(u)(y)$.

Network architecture. To accommodate the probabilistic prediction for DeepONets, we split the final layers of both the Branch Net and the Trunk as depicted in Figure 5. These new layers provide independent processing for the mean μ_θ and the standard deviation σ_θ of the DeepONet's Gaussian estimates of the output. To avoid numerical instability, we follow [29] and construct the DeepONet to predict the log standard deviation $\log \sigma_\theta(u)(y)$ and recover the standard deviation using the exponential function, *i.e.*, $\sigma_\theta(u)(y) = \exp(\log \sigma_\theta(u)(y))$.

Probabilistic Training. During training, our goal is to calibrate the DeepONet's parameters θ to produce the Gaussian estimates of predictive uncertainty. This requires designing a probabilistic training protocol to maximize the associated likelihood of observing the target post-fault




Figure 5: The Probabilistic DeepONet framework. The DeepONet outputs are interpreted as parameters of a Gaussian distribution $\mathcal{N}(\mu_\theta(u)(y), \sigma_\theta^2(u)(y))$ of the *true* values $G^\dagger(u)(y)$. The loss function is given by the negative log-likelihood (11).

trajectory values. Equivalently, one can minimize the negative-log likelihood of the target post fault trajectories, *i.e.*,

$$\mathcal{L}_p(\theta) = \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \frac{(\mu_\theta(u_i)(y_i) - G^\dagger(u_i)(y_i))^2}{\sigma_\theta^2(u_i)(y_i)} + \frac{1}{2} \log(2\pi\sigma_\theta^2(u_i)(y_i)). \quad (11)$$

As described in [29], during probabilistic training, this negative-log likelihood loss guides the DeepONet to build trustworthiness gradually by decreasing the predicted standard deviation σ_θ .

Trustworthy prediction of post-fault trajectories. To predict and estimate the aleatoric uncertainty of a post fault trajectory $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$, Prob-DeepONet predicts the mean trajectory $\{\mu_\theta(u)(y) : y \in Y_{\text{mesh}}\}$ and the standard deviation $\{\sigma_\theta(u)(y) : y \in Y_{\text{mesh}}\}$. The mean trajectory represents the prediction of the true trajectory $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$ and we use the standard deviation to construct an estimate of its uncertainty. Let us conclude this section by noting that compared to B-DeepONet, which requires M DeepONet forward passes to construct an estimate of the uncertainty (see Figure 4), Prob-DeepONet estimates uncertainty instantaneously at virtually no extra cost, *i.e.*, it uses only *one* Prob-DeepONet forward pass. In the next section, we use a series of numerical experiments to demonstrate the B-DeepONet and Prob-DeepONet’s effectiveness in providing trustworthy predictions.

5 Numerical Experiments

This section presents four numerical experiments aimed at demonstrating the predictive power and reliability of the proposed B-DeepONet and Prob-DeepONet. The first experiment (Section 5.3) compares the predictive power of B-DeepONet, Prob-DeepONet, and vanilla DeepONet. The second experiment (Section 5.4) builds confidence intervals to verify how reliable the predictions of B-DeepONet and Prob-DeepONet are. The third experiment (Section 5.5) tests the performance of the proposed methods when the input fault-on trajectory is contaminated with noise. The final experiment (Section 5.6) tests how well the proposed methods predict violations of pre-defined voltage limits. We start this section by describing the power grid simulation framework used to generate the datasets.




Figure 6: One-line diagram of the New York-New England 16-generator 68-bus power grid.

5.1 Dataset generation

We used time-domain simulations on the New York-New England 16-generator 68-bus power grid model (see Figure 6) to generate the training and testing datasets. We performed the simulations using EPTOOL –a tool developed based on the Power System Toolbox [8]. We also considered the full dynamical model of the generators, including turbine governors, excitation systems, and power system stabilizer.

In the numerical experiments, we learn the operator $G^\dagger$ that predicts the post-fault trajectories of the voltage magnitude of bus $h = 19$ (see Figure 6). However, we remark that the proposed framework can be implemented on other buses or states, *e.g.*, frequency. We leave for our future work the design of a transfer learning strategy to reduce the training time of DeepONets for other buses and states (*e.g.*, the frequency at some bus j) using a trained DeepONet of a selected state (*e.g.*, the voltage of bus h , where $h \neq j$).

To train and test the B-DeepONet and Prob-DeepONet, we generate two data pools of transient response trajectories. We generate the first data pool of 1000 trajectories by simulating the grid’s transient response to $N-1$ faults. Similarly, we generate the second data pool of 1000 trajectories by simulating $N-2$ faults. Each trajectory is simulated for $T = 9$ seconds with a resolution of 100 [Hz]. As a result, all the $N = 2000$ trajectories have $N_{\text{samples}} = 900$ samples.

To simulate faults, we proceed as follows. We assume the power grid operates on a stable equilibrium before the fault, *i.e.*, within the interval $[0, t_f)$. At the time t_f , the grid experiences the fault. We simulate $N-1$ (resp. $N-2$) faults by disconnecting uniformly at random one (resp. two) transmission lines. Then, at time t_{cl} , we assume the reclosing system clears the fault. Finally, the post-fault trajectory occurs within the interval $(t_{cl}, T]$.

To simulate each fault, we must define the values for t_f and t_{cl} . Here, for all N trajectories, we fix the value of the clearing time to $t_{cl} = 2$ seconds. Let Δt_f denote the duration time of the fault, *i.e.*, $\Delta t_f := t_{cl} - t_f$. We simulate N Δt_f ’s from the uniform distribution $\mathcal{U}(0.2, 0.5)$. Then, for each simulated Δt_f , we obtain t_f via $t_f = 2.0 - \Delta t_f$.

Finally, to generate the training and testing data pools, we combine and shuffle the $N-1$ and $N-2$ data pools and perform a 70 – 30% split.

Training dataset $\mathcal{D}_{\text{train}}$. Here, we use the training data pool of N_{train} transient voltage trajectories to construct our training dataset $\mathcal{D}_{\text{train}}$. We first split the transient simulation domain $[0, T]$

into the two adjacent domains. $T_u = [0, t_{cl}]$ and $T_G = (t_{cl}, T]$. T_u contains the fault-on information and T_G represents the domain of post-fault trajectories.

The Branch Net u takes as inputs the voltage trajectories in the interval T_u . We discretize these inputs using $m = 200$ sensors. Thus, the discretized input $(u(x_1), \dots, u(x_m))$ coincides with the voltage trajectory simulated at a resolution of 100 [Hz]. One can also use a different number of sensors. However, this requires performing interpolation.

On the other hand, the Trunk Net takes as inputs time values y sampled uniformly from the interval T_G . Finally, the target operator values $G^\dagger(u)(y)$ for supervised training correspond to the transient voltage response evaluated at time y for the given input (fault-on trajectory) u . Note that if y does not coincide with the simulation resolution of 100 [Hz], we need to perform interpolation. The training dataset is then

$$\mathcal{D}_{\text{train}} = \{(u_i(x_1), \dots, u_i(x_m)), y_i, G^\dagger(u_i)(y_i)\}_{i=1}^{N_{\text{train}}}.$$

In the above $\mathcal{D}$, we have shown the case when only $Q = 1$ point $y_i \in T_G$ is sampled per input trajectory u_i . However, we remark that we can always sample multiple points $\{y_i^j\}_{j=1}^Q$ with $Q > 1$ per input trajectory u_i . In all our experiments, we sampled $Q = 10$ points.

Testing dataset $\mathcal{D}_{\text{test}}$. The testing dataset consists of all the N_{test} simulated trajectories of the testing data pool. Using the DeepONet notation, the post-fault component of a testing trajectory is denoted as $\{G(u)(y) : y \in Y_{\text{mesh}}\}$. This post-fault component corresponds to the transient trajectory evaluated at points $y \in Y_{\text{mesh}} \subset T_G$ for the given input u . In all our experiments, we selected Y_{mesh} to be the set of 500 linearly spaced points within the post-fault interval T_G . Finally, to evaluate the performance of B-DeepONet and Prob-DeepONet on the testing dataset, we use two metrics: the L_1 and the L_2 relative errors.

5.2 Training protocols and neural networks

Training protocols. We implemented Prob-DeepONet and B-DeepONet using PyTorch and published the code in GitHub.³ To train the Prob-DeepONet and vanilla DeepONet, we used the Adam [16] optimizer with default hyper-parameters and initial learning rate $\eta = 10^{-4}$. We reduced this learning rate whenever the loss reached a plateau or increased. For B-DeepONet, we manually calibrated all the parameters of SGHMC (Algorithm 1).

Neural networks. We implemented the Branch and the Trunk Nets using the modified fully-connected architecture proposed in [27] and used in our previous paper [20]. The forward pass of this modified network is

$$\begin{aligned} U &= \phi(XW^1 + b^1), \quad V = \phi(XW^2 + b^2) \\ H^{(1)} &= \phi(XW^{z,1} + b^{z,1}) \\ Z^{(k)} &= \phi(H^{(k)}W^{z,k} + b^{z,k}), \quad k = 1, \dots, d \\ H^{(k+1)} &= (1 - Z^{(k)}) \odot U + Z^{(k)} \odot V, \quad k = 1, \dots, d \\ f_\theta(x) &= H^{(d+1)}W + b, \end{aligned}$$

Here, X is the input tensor to the neural network, d is the number of hidden layers (*i.e.*, the network's depth), $\odot$ is the element-wise product, and ϕ , in this paper, is a point-wise sinusoidal activation function. The trainable parameters are

$$\theta = \{W^1, b^1, W^2, b^2, \{W^{z,l}, b^{z,l}\}_{l=1}^d, W, b\}.$$

We initialized all these parameters using the Glorot uniform algorithm.

³The code will be available after publication.

5.3 Experiment 1: Mean prediction and generalization

This experiment compares the mean predictive power and the generalization capability of B-DeepONet, Prob-DeepONet, and vanilla DeepONet. To this end, we train the models for 10000 epochs. Once we complete training, we select the best-trained models for vanilla DeepONet and Prob-DeepONet, and the M -ensemble produced by B-DeepONet. We use these trained models to compute the average and standard deviation (st.dev.) of the L_1 - and L_2 - relative errors for 100 test trajectories selected uniformly at random from $\mathcal{D}_{\text{test}}$.

We report the results in Table 1. These results show that B-DeepONet and Prob-DeepONet provide similar mean predictive power and generalization capability to the vanilla DeepONet. We believe this is a remarkable result because our proposed methods were not designed to provide improved mean prediction of post-fault trajectories.

	B-DeepONet	Prob-DeepONet	vanilla DeepONet
mean L_1	1.81 %	1.62 %	1.63 %
st.dev. L_1	1.88 %	2.02 %	2.0 %
mean L_2	2.23 %	1.95 %	2.08 %
st.dev. L_2	2.40 %	2.51 %	2.57 %

Table 1: The average and standard deviation (st.dev.) of the L_1 - and L_2 - relative errors of 100 trajectories selected uniformly at random from $\mathcal{D}_{\text{test}}$.

5.4 Experiment 2: Reliability

This experiment verifies how well the M-ensemble produced by B-DeepONet and the trained Prob-DeepONet quantify the uncertainty. To this end, we construct a 95% confidence interval (CI) for post-fault trajectory prediction. B-DeepONet constructs the confidence interval by computing the standard deviation over the M simulated post-fault trajectories $\{G_{\theta_{N+1-k}}(u)(y) : y \in Y_{\text{mesh}}\}_{k=1}^M$. Prob-DeepONet constructs the confidence interval using the point-wise predicted standard deviation of the post fault trajectory, that is, $\{\sigma_{\theta^*}(u)(y) : y \in Y_{\text{mesh}}\}$.

To measure how reliable is the predicted uncertainty, we define the following ratio ϵ_{ratio} :

$$\epsilon_{\text{ratio}} = \frac{\# \text{ of points of the true post-fault trajectory } \{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\} \text{ in the CI}}{\# \text{ of points of the true post-fault trajectory } \{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}} \cdot 100\%$$

Note that ϵ_{ratio} measures the percentage of the true post-fault test trajectory captured by the predicted confidence interval.

Figure 7 depicts the mean prediction and confidence intervals generated by the proposed methods for four test trajectories selected uniformly at random from $\mathcal{D}_{\text{test}}$. One should note that the constructed CIs capture pretty well the true test post-fault trajectories, *i.e.*, their ratio is $\epsilon_{\text{ratio}} = 100\%$.

We also report in Table 2 the average ratio ϵ_{ratio} for 100 trajectories selected from $\mathcal{D}_{\text{test}}$. As expected, both B-DeepONet and Prob-DeepONet capture with their confidence intervals $\approx 95\%$ of all the test trajectories.

	B-DeepONet	Prob-DeepONet
ϵ_{ratio}	93.32 %	94.27 %

Table 2: The average ϵ_{ratio} of all the post-fault trajectories in $\mathcal{D}_{\text{test}}$.




Figure 7: *Top:* B-DeepONet mean prediction and confidence interval (CI) for two post-fault trajectories selected at random from $\mathcal{D}_{\text{test}}$. The CI contains $\epsilon_{\text{ratio}} = 100\%$ of the true test trajectories. *Bottom:* Prob-DeepONet mean prediction and confidence interval for two post-fault trajectories selected at random from $\mathcal{D}_{\text{test}}$. The CI contains $\epsilon_{\text{ratio}} = 100\%$ of the true test trajectories.

To conclude this experiment, we compute, for B-DeepONet and Prob-DeepONet, the fraction of points of the test trajectories falling within a specified standard deviation range. To this end, let us first define for B-DeepONets the predicted mean of the test trajectory $\{G^\dagger(u)(y) : y \in Y_{\text{mesh}}\}$ as $\{\hat{\mu}_\theta^B(u)(y) : y \in Y_{\text{mesh}}\}$ and the predicted standard deviation as $\{\hat{\sigma}_\theta^B(u)(y) : y \in Y_{\text{mesh}}\}$.

Formally, we aim to compute the average number of points where the true post-fault trajectory differs from the predictive mean by less than a factor χ of the predicted standard deviation:

$$\frac{1}{|\mathcal{D}_{\text{test}}|} \frac{1}{|Y_{\text{mesh}}|} \sum_{u \in \mathcal{D}_{\text{test}}} \sum_{y \in Y_{\text{mesh}}} \mathbb{1}\{E(u)(y)\}, \quad (12)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function and $E(u)(y)$ is the aforementioned event:

$$E(u)(y) := |\hat{\mu}(u)(y) - G^\dagger(u)(y)| \leq \chi \cdot \hat{\sigma}(u)(y).$$

Here, $\hat{\mu} \equiv \hat{\mu}_\theta^B$ and $\hat{\sigma} \equiv \hat{\sigma}_\theta^B$ for B-DeepONets, and $\hat{\mu} \equiv \mu_\theta$ and $\hat{\sigma} \equiv \sigma_\theta$ for Prob-DeepONet.

Figure 8 compares the average points obtained using (12) for B-DeepONet and Prob-DeepONet against $\mathbb{P}(|\mathcal{N}(0,1)| \leq \chi)$, *i.e.*, the average number of points of a Normal distribution lying within a factor χ of the unit standard deviation. As expected, the Prob-DeepONet result remains close to the Normal Distribution. Such a result confirms our assumption of Gaussian distributed network errors. Remarkably, the B-DeepONet average remains also close to the Normal distribution average. This experimental evidence suggests that the proposed methods adequately measure the predictive uncertainty.




Figure 8: Quantitative analysis of the B-DeepONet (*left*) and Prob-DeepONet (*right*) predicted standard deviation values for all post-fault trajectories in $\mathcal{D}_{\text{test}}$. We show the percentage of points where the difference between the mean prediction and true solution is less than a factor χ of the predicted standard deviation. They both closely follow the cumulative distribution of the half-normal distribution.

5.5 Experiment 3: Robustness to noise

In this experiment, we verify the performance of the proposed methods when the test input u , containing fault-on information, is contaminated with noise. This scenario can happen in the online setting, which uses noisy measurements collected using PMUs.

We assume the noisy discretized inputs $\tilde{u}(x_i)$, for $i = 1, \dots, m$ are independently Gaussian distributed centered at the true input value, *i.e.*,

$$\tilde{u}(x_i) = u(x_i) + \epsilon_i, \quad i = 1, \dots, m,$$

where ϵ_i is an independent Gaussian noise with zero mean and standard deviation σ , *i.e.*, $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. We select the noise level to be $\sigma = 0.01$.

Figure 9 depicts the predicted confidence intervals for four test trajectories. The results show that the uncertainty estimate, represented by the 95% confidence interval, does not deteriorate in this noisy scenario.

We conclude this experiment by showing that B-DeepONet and Prob-DeepONet are robust to noisy inputs. To this end, we report in Table 3 the ϵ_{ratio} for 100 test trajectories selected from $\mathcal{D}_{\text{test}}$ whose inputs are contaminated with noise. The results show no deterioration; the 95% confidence intervals still capture about 94% of the true post-fault trajectories.

	B-DeepONet	Prob-DeepONet
ϵ_{ratio}	92.87 %	93.27 %

Table 3: The average ϵ_{ratio} for 100 post-fault trajectories selected from $\mathcal{D}_{\text{test}}$ whose inputs are contaminated with noise.

5.6 Experiment 4: A power grid application

This final experiment verifies how well the proposed B-DeepONet and Prob-DeepONet predict violations of pre-established under-voltage profiles. Such a violation could trigger, for example, under-voltage load shedding. We refer the interested reader to [30, 21] and references therein for more information about these voltage profiles.




Figure 9: *Top:* B-DeepONet mean prediction and confidence interval (CI) for two post-fault trajectories with inputs contaminated with noise and selected at random from $\mathcal{D}_{\text{test}}$. The CI contains $\epsilon_{\text{ratio}} = 100\%$ of the true test trajectories. *Bottom:* Prob-DeepONet mean prediction and confidence interval for two post-fault trajectories with inputs contaminated with noise and selected at random from $\mathcal{D}_{\text{test}}$. The CI in this noisy scenario contains $\epsilon_{\text{ratio}} = 100\%$ of the true test trajectories.

We illustrate in Figure 10 a test trajectory for B-DeepONet and Prob-DeepONet and their respective voltage profiles. One can observe that the 95% confidence interval predicted by B-DeepONet and Prob-DeepONet may intersect the voltage profiles at some point. Furthermore, we present in Figure 11 a regression scatter plot of the mean predictions and 95% confidence intervals for all test trajectories at time $y = 2.2$ seconds. The figure compares the mean prediction with the true value at time $y = 2.2$ seconds. It also verifies whether the CIs intersect regions that can trigger under-voltage load shedding. In particular, we say a false-negative (FN) alarm occurs when the whole CI lies within the FN region. For a false-positive (FP) alarm, we consider two scenarios. 1) In the conservative scenario, we say an FP occurs when the CI intersects the FP region. 2) In a non-conservative scenario, the FP occurs when the whole CI lies within the FP region.

The results of Figure 11 illustrate that the proposed methods eliminate all false-negative alarms. Moreover, the results show that the methods yield a 32% for B-DeepONet and 29% for Prob-DeepONet (resp. 0%) false-positive rate in the conservative (resp. non-conservative) scenario.

6 Discussion and Future Work

We start this section by providing a summary of our results.

Summary of our results. Our first main result illustrates (in Figure 1) that B-DeepONet and




Figure 10: The 95% confidence interval predicted by the B-DeepONet (*left*) and Prob-DeepONet (*right*) compared with the profile for under-voltage load shedding.




Figure 11: A regression scatter plot of the B-DeepONet (*left*) and Prob-DeepONet (*right*) mean predictions (dot) and 95% confidence intervals (error bar) for all test trajectories at time $y = 2.2$ seconds. We verify whether the predicted CIs intersect regions that can trigger under-voltage load shedding. .

Prob-DeepONet provide comparable mean prediction and generalization than the vanilla DeepONet trained with Adam. Our second main result validates (in Figure 7 and Table 2) that B-DeepONet and Prob-DeepONet effectively quantify the aleatoric uncertainty, and, hence, they can provide reliable post-fault trajectory predictions. In particular, the predicted confidence intervals capture almost 95% of the true test trajectories. We achieve these results using a small training dataset. Our final result illustrates (in Figure 10 and Figure 11) that the reliable predictions of the proposed methods enable reducing (to almost zero) the number of false-negative alarms. Such alarms are raised when the methods wrongly predict that a post-fault voltage trajectory violates pre-defined limits.

Our future work seeks to improve the design of DeepONets for transient stability prediction as follows.

On relaxing the DeepONet Sensors. We note that in practice, the set of sensor evaluation points for DeepONets ($x_1, \dots, x_m$) (which corresponds to the set of sampling times of a transient response simulator or PMUs) can be affected by delays/noise or not known a-priori. Thus, it is part of our future work to relax the sensor location constraint by designing a resolution-independent DeepONet.

On using neighboring information. In this paper, we did not use trajectory information from neighboring buses to train our DeepONets. In our previous work [18], we observed that using

this neighboring information improves training and generalization. Thus, it is part of our future work to design a multi-input multi-output DeepONet.

On developing a transfer learning strategy. We plan to design a transfer learning strategy that enables using a trained B-DeepONet or Prob-DeepONet to train efficiently multiple post-fault trajectory predictors on a large-scale power grid.

On power-grid informed DeepONets. Physics-Informed Neural Networks (PINN) [24] are neural networks that we train to satisfy the underlying mathematical physics of the given problem. PINNs have provided us with multiple transformative results by using the physics of the underlying problem as a regularizing agent. This regularization limits the space of possible solutions, which, in turn, enables generalizing well even when the amount of training data is small. In our previous work [20], we developed a PINN to simulate stiff and nonlinear differential-algebraic equations. In our ongoing work, we are designing DeepONets that use the power grid dynamics as a regularizing agent. We aim to use these power grid-informed DeepONets to enhance generalization capabilities for transient analysis and build a transient analysis tool for large-scale power grids.

7 Conclusion

This paper proposes using Deep Operator Networks (DeepONet) to predict post-fault trajectories reliably. In particular, we design a DeepONet to (1) take as inputs the fault-on trajectories collected, for example, via simulation or phasor measurement units, and (2) provide as outputs the predicted post-fault trajectories. Furthermore, we design two uncertainty quantification methods that provide the much-needed ability to balance efficiency with reliable/trustworthy predictions. First, we design a *Bayesian DeepONet* (B-DeepONet) that uses stochastic gradient Hamiltonian Monte-Carlo to sample from the posterior distribution of the DeepONet parameters. Second, we design a *Probabilistic DeepONet* (Prob-DeepONet) that uses a probabilistic training strategy to equip DeepONets with a form of automated uncertainty quantification, at virtually no extra computational cost. Finally, we use a series of experiments on the IEEE-68-bus system that provide the empirical evidence about the effectiveness of the proposed B-DeepONet and Prob-DeepONet.

Acknowledgment

The authors gratefully acknowledge the support of the National Science Foundation (DMS-1555072, DMS-1736364, DMS-2053746, and DMS-2134209), and Brookhaven National Laboratory Subcontract 382247, and U.S. Department of Energy (DOE) Office of Science Advanced Scientific Computing Research program DE-SC0021142.

References

- [1] F. Alvarado and S. Oren. Transmission system operation and interconnection. *National transmission grid study—Issue papers*, pages A1–A35, 2002.
- [2] M. Anghel, J. Anderson, and A. Papachristodoulou. Stability analysis of power systems using network decomposition and local gain analysis. pages 1–7, 2013.

- [3] S. Y. Caliskan and P. Tabuada. Compositional transient stability analysis of multimachine power networks. *IEEE Transactions on Control of Network systems*, 1(1):4–14, 2014.
- [4] T. Chen and H. Chen. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE Transactions on Neural Networks*, 6(4):911–917, 1995.
- [5] T. Chen, E. Fox, and C. Guestrin. Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning*, pages 1683–1691. PMLR, 2014.
- [6] H.-D. Chiang. *Direct methods for stability analysis of electric power systems: theoretical foundation, BCU methodologies, and applications*. John Wiley & Sons, 2011.
- [7] H.-D. Chiang, F. F. Wu, and P. P. Varaiya. A bcu method for direct analysis of power system transient stability. *IEEE Transactions on Power Systems*, 9(3):1194–1208, 1994.
- [8] J. H. Chow and K. W. Cheung. A toolbox for power system dynamics and control engineering education and research. *IEEE transactions on Power Systems*, 7(4):1559–1564, 1992.
- [9] F. Dorfler and F. Bullo. Synchronization and transient stability in power networks and nonuniform kuramoto oscillators. *SIAM Journal on Control and Optimization*, 50(3):1616–1642, 2012.
- [10] A. Gupta, G. Gurralla, and P. Sastry. An online power system stability monitoring system using convolutional neural networks. *IEEE Transactions on Power Systems*, 34(2):864–872, 2018.
- [11] M. He, J. Zhang, and V. Vittal. Robust online dynamic security assessment using adaptive ensemble decision-tree learning. *IEEE Transactions on Power systems*, 28(4):4089–4098, 2013.
- [12] I. A. Hiskens and D. J. Hill. Energy functions, transient stability and voltage behaviour in power systems with nonlinear loads. *IEEE transactions on power systems*, 4(4):1525–1533, 1989.
- [13] A. Iserles. *A first course in the numerical analysis of differential equations*. Number 44. Cambridge university press, 2009.
- [14] J. James, D. J. Hill, A. Y. Lam, J. Gu, and V. O. Li. Intelligent time-adaptive transient stability assessment system. *IEEE Transactions on Power Systems*, 33(1):1049–1058, 2017.
- [15] L. Jin, H. Liu, R. Kumar, J. D. Mc Calley, N. Elia, and V. Ajjarapu. Power system transient stability design using reachability based stability-region computation. pages 338–343, 2005.
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [17] P. Kundur. Power system stability. *Power system stability and control*, pages 7–1, 2007.
- [18] J. Li, M. Yue, Y. Zhao, and G. Lin. Machine-learning-based online transient analysis via iterative computation of generator dynamics. In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGrid-Comm)*, pages 1–6. IEEE, 2020.

- [19] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- [20] C. Moya and G. Lin. Dae-pinn: A physics-informed neural network model for simulating differential-algebraic equations with application to power networks. *arXiv preprint arXiv:2109.04304*, 2021.
- [21] C. Mozina. Undervoltage load shedding. In *2007 Power Systems Conference: Advanced Metering, Protection, Control, Communication, and Distributed Resources*, pages 39–54. IEEE, 2007.
- [22] R. M. Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- [23] M. Pai, K. Padiyar, and C. Radhakrishna. Transient stability analysis of multi-machine ac/dc power systems via energy-function method. *IEEE Transactions on Power Apparatus and Systems*, (12):5027–5035, 1981.
- [24] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [25] R. Schainker, P. Miller, W. Dubbelday, P. Hirsch, and G. Zhang. Real-time dynamic security assessment: fast simulation and modeling applied to emergency outage security of the electric grid. *IEEE Power and Energy magazine*, 4(2):51–58, 2006.
- [26] P. Varaiya, F. F. Wu, and R.-L. Chen. Direct methods for transient stability analysis of power systems: Recent results. *Proceedings of the IEEE*, 73(12):1703–1715, 1985.
- [27] S. Wang, Y. Teng, and P. Perdikaris. Understanding and mitigating gradient pathologies in physics-informed neural networks. *arXiv preprint arXiv:2001.04536*, 2020.
- [28] L. Wehenkel, T. Van Cutsem, and M. Ribbens-Pavella. An artificial intelligence framework for online transient stability assessment of power systems. *IEEE Transactions on Power Systems*, 4(2):789–800, 1989.
- [29] N. Winovich, K. Ramani, and G. Lin. Convnpde-uq: Convolutional neural networks with quantified uncertainty for heterogeneous elliptic partial differential equations on varied domains. *Journal of Computational Physics*, 394:263–279, 2019.
- [30] Y. Zhang, M. Yue, and J. Wang. Adaptive load shedding for grid emergency control via deep reinforcement learning. *arXiv preprint arXiv:2102.12908*, 2021.
- [31] L. Zhu, D. J. Hill, and C. Lu. Hierarchical deep learning machine for power system online transient stability prediction. *IEEE Transactions on Power Systems*, 35(3):2399–2411, 2019.

Appendix 1

In this Appendix, we empirically show that the network errors ($G_\theta(u)(y) - G^\dagger(u)(y)$) the vanilla DeepONet experiences, during gradient descent training (*e.g.*, Adam), follows a Gaussian distribution. To this end, we plot the network errors and their distributions (see Figure 12) after




Figure 12: *Left:* network errors experienced by the vanilla DeepONet trained with Adam. *Right:* distribution of the network errors. *Top:* network errors and their distribution when $\mathcal{D}_{\text{train}}$ is simulated using single-line contingencies. *Middle:* network errors and their distribution when $\mathcal{D}_{\text{train}}$ is simulated using double-line contingencies. *Bottom:* network errors and their distribution when $\mathcal{D}_{\text{train}}$ is simulated using both single-line and double-line contingencies.

some burn-in number (80000) of steps⁴. In Figure 12, we have considered three different cases. In the first case, the training data $\mathcal{D}_{\text{train}}$ contains voltage fault-on and post-fault trajectories simulated from single-line contingencies. In the second case, $\mathcal{D}_{\text{train}}$ contains voltage trajectories simulated from double-line contingencies. In the last case, $\mathcal{D}_{\text{train}}$ contains voltage trajectories simulated from both single-line and double-line contingencies. The results show that regardless of the case, the network errors experienced by DeepONets are Gaussian distributed. Such a result provides us with empirical confirmation that the outputs of the Prob-DeepONet can be modeled as the parameters of a Gaussian distribution $\mathcal{N}(\mu_{\theta}(u)(y), \sigma_{\theta}^2(u)(y))$.

⁴One should note that a step is different from an epoch. Here we define *step* as one gradient descent execution.