Downloaded from orbit.dtu.dk on: Apr 26, 2024

DTU Library

Physics-Informed Neural Networks for Power Systems

Misyris, Georgios; Venzke, Andreas; Chatzivasileiadis, Spyros

Published in:
Proceedings of 2020 IEEE Power & Energy Society General Meeting

Publication date:
2020

Document Version
Peer reviewed version

Link back to DTU Orbit

Citation (APA):
Misyris, G., Venzke, A., & Chatzivasileiadis, S. (2020). Physics-Informed Neural Networks for Power Systems. In *Proceedings of 2020 IEEE Power & Energy Society General Meeting* IEEE.

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Physics-Informed Neural Networks for Power Systems

George S. Misyris, Andreas Venzke and Spyros Chatzivasileiadis
Center for Electric Power and Energy
Technical University of Denmark
{gmisy, andven, spchatz}@elektro.dtu.dk

arXiv:1911.03737v3 [eess.SY] 29 Jan 2020

***Abstract*—This paper introduces for the first time, to our knowledge, a framework for physics-informed neural networks in power system applications. Exploiting the underlying physical laws governing power systems, and inspired by recent developments in the field of machine learning, this paper proposes a neural network training procedure that can make use of the wide range of mathematical models describing power system behavior, both in steady-state and in dynamics. Physics-informed neural networks require substantially less training data and can result in simpler neural network structures, while achieving high accuracy. This work unlocks a range of opportunities in power systems, being able to determine dynamic states, such as rotor angles and frequency, and uncertain parameters such as inertia and damping at a fraction of the computational time required by conventional methods. This paper focuses on introducing the framework and showcases its potential using a single-machine infinite bus system as a guiding example. Physics-informed neural networks are shown to accurately determine rotor angle and frequency up to *87 times faster* than conventional methods.**

***Index Terms*—deep learning, neural network, power system dynamics, power flow, system inertia.**

I. Introduction

Machine learning techniques demonstrate impressive results for a range of highly complex tasks, especially where an accurate mathematical representation of the problem cannot be obtained. Applications include image recognition, robotics, weather forecasting, and others [1]. In power systems, decision trees and neural networks have been shown to solve computational problems both in dynamics and optimization at a fraction of the time required by traditional approaches, being up to three order of magnitude faster [2]–[6].

Up to this point, however, machine learning methods applied to power systems (and other physical systems) were largely agnostic to the underlying physical model. This made them heavily dependent on the quality of the training data, it required large training datasets, and oftentimes complex neural network structures. Despite recent efforts for efficient creation of datasets with encouraging results [7], [8], generating the required training dataset size still requires substantial computational effort. In this work, inspired by [9], [10], we reduce the dependency on training data and complex neural network structures by exploiting *inside the neural network training* the underlying physical laws described by power system models.

This is the first work, to our knowledge, that proposes physics-informed neural networks for power system applications. It introduces a neural network training framework that can exploit the underlying physical laws and the available power system models both for steady-state and dynamics. Following recent approaches reported in [9], [10], we incorporate the power system differential and algebraic equations inside the training procedure. Exploiting advances in automatic differentiation [11] that are implemented in Tensorflow [12], we can directly compute derivatives of neural network outputs during training, such as the rotor angle, and build neural networks able to accurately capture the rotor angle and frequency dynamics. Our approach (i) requires less initial training data, (ii) can result to smaller neural networks, while (iii) demonstrating high performance.

Physics-informed neural networks introduce a novel technology that may lead to a new class of numerical solvers [10] as well as dynamic state estimation techniques [13]. Within power systems, they have the potential to solve systems of differential-algebraic equations at a fraction of computational time required for conventional methods, are able to directly determine the value of state variables at any time instant t_1 (without the need to integrate from t_0 to t_1), and can handle directly higher-order differential equations without the need to introduce additional variables to solve a first-order system.

In this paper, we present the main principles for the application of physics-informed neural networks of [10] in power systems, focusing on power system dynamics and using the swing equation as an example. Besides obtaining solutions to ordinary differential equations, we demonstrate how the same methods can be used to estimate uncertain parameters such as inertia and damping. The contributions of our work are:

1) We propose physics-informed neural networks to (i) accurately determine solutions of differential equations and, thus, values of power system dynamic states, such as rotor angle and frequency, and (ii) identify uncertain power system parameters. Contrary to previous approaches, physics-informed neural networks utilize the underlying physical model, lead to significantly reduced computation time and need less training data.
2) For the single machine infinite bus (SMIB) system, we show that physics-informed neural networks (i) predict system dynamics with high accuracy at a fraction of the computational time required by conventional approaches (28-87 times faster in our study), and (ii) can identify with high accuracy uncertain system parameters such as inertia and damping.

This paper is structured as follows: Section II describes the employed power system model and introduces the architecture of physics-informed neural networks. Section III presents simulation results demonstrating the performance of physics-informed neural networks. Section IV discusses the challenges and the opportunities emerging from the successful application of this concept. Section V concludes. The code to reproduce

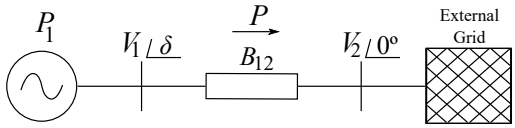

Fig. 1. Single machine infinite bus system

the simulation results is available online [14].

II. Methodology

A. Physical Model for Power System Dynamics

Power system dynamics, in their simplest and most common form, are described by the swing equation, neglecting transmission losses and bus voltage deviations. For each generator k, the resulting system of equations can then be represented by [15], [16]:

$$m_k\ddot{\delta}_k + d_k\dot{\delta}_k + \sum_j B_{kj}V_kV_j\sin(\delta_k - \delta_j) - P_k = 0 \quad (1)$$

where m_k defines the generator inertia constant, d_k represents the damping coefficient, B_{kj} is the $\{k, j\}$-entry of the bus susceptance matrix, P_k is the mechanical power of the kth generator, V_k, V_j are the voltage magnitudes at buses k, j and δ_k, δ_j represent the voltage angles behind the transient reactance. $\dot{\delta}_k$ is the angular frequency of generator k, often also denoted as ω_k.

1) Single Machine Infinite Bus (SMIB) System: The single-machine infinite-bus system, shown in Fig. 1, has been widely used to understand and analyze the fundamental dynamic phenomena occurring in power systems. As the focus of this paper is on the introduction of physics-informed neural networks for power systems, we will use this system as a guiding example. Note though that our proposed framework is general. Future work will focus on larger, more complex systems. The swing equation (1) for the SMIB system is given by:

$$m_1\ddot{\delta} + d_1\dot{\delta} + B_{12}V_1V_2\sin(\delta) - P_1 = 0 \quad (2)$$

In the rest of this paper, we will show how physics-informed neural networks can accurately estimate both rotor angle δ and frequency $\dot{\delta}$ while P_1 varies within $[P_{\min}, P_{\max}]$, and can identify uncertain parameters such as m_1 and d_1.

B. Physics-Informed Neural Networks

In the following, we explain the general architecture of physics-informed neural networks, and detail its application to the SMIB system. Feed-forward neural networks are composed of the input layer, fully connected hidden layers having a non-linear activation function at each neuron, and the output layer. Between each layer a weight matrix $\mathbf{W}$ and bias $\mathbf{b}$ is applied. During training, weight matrices and biases are optimized to minimize an objective function which usually penalizes the deviation of the neural network prediction from the training data. Neural networks are universal function approximators as they can, in theory, learn any unknown function between some inputs and outputs. Therefore, neural networks could be used to directly learn the nonlinear mapping between the inputs and the outputs of differential equations, such as (2). Not taking into account the underlying physical model, however, will require large amounts of training data and a large neural network size. The work in [10] introduced a framework for physics-informed neural networks which we will rely on in the following. Considering physical laws during training allows to bound the space of admissible solutions to the neural network parameters, which translates to a lower requirement in both the amount of training data and neural network size.

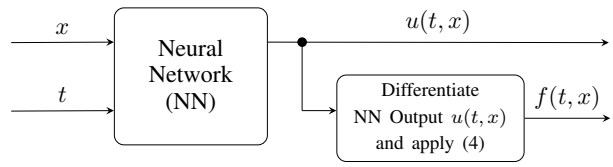

Fig. 2. General structure of a physics-informed neural network: it predicts the output $u(t, x)$ given inputs x and t. Then, using automatic differentiation [11] of the same neural network, the partial derivatives of $u(t, x)$ are computed, and $f(t, x)$ is evaluated. The parameters λ are either assumed to be known, or are optimized as part of the neural network training. During training, the neural network weights and biases are adjusted according to loss function (5), which minimizes the deviation of both the output prediction $u(t, x)$ from ground truth and $f(t, x)$ from 0.

Following notation similar to Ref. [10], the general form of the functions that the physics-informed neural network can approximate is:

$$\frac{\partial u}{\partial t} = -N[u;\lambda]\,,\ x \in \Omega,\ t \in [0, T] \quad (3)$$

where $u(t, x)$ is the solution and $N[u; \lambda]$ is a nonlinear operator connecting the state variables u with the system parameters λ. The term t denotes time and x the system input. The domain Ω can be bounded based on prior knowledge of the dynamical system and $[0, T]$ is the time interval within which the system evolves. The model parameters λ can be constant or unknown. In case λ is unknown, the problem of approximating function (3) becomes a problem of system identification, where we seek parameters λ for which the expression in (3) is satisfied. To enforce the physical law describing the dynamical system we define the physics-informed neural network $f(t, x)$:

$$f(t,x) = \frac{\partial u}{\partial t} + N[u,\lambda] \quad (4)$$

Note that if the system parameters λ are known the nonlinear operator $N[u, \lambda]$ simplifies to $N[u]$. The overall architecture is shown in Fig. 2. A neural network is used to predict $u(t, x)$ based on the inputs t and x. To determine $f(t, x)$, we use automatic differentiation [11] of the components of the neural network predicting $u(t, x)$. Based on this, we compute the required derivatives of $u(t, x)$ with respect to time t and system inputs x. As a result, the neural network predicting $f(t, x)$ has the same parameters compared to the neural network predicting $u(t, x)$, but different activation functions. The shared parameters of the two neural networks are optimized by minimizing the loss function:

$$MSE = \underbrace{\tfrac{1}{N_u}\sum_i^{N_u}|u(t_u^i, x_u^i) - u^i|^2}_{MSE_u} + \underbrace{\tfrac{1}{N_f}\sum_i^{N_f}|f(t_f^i, x_f^i)|^2}_{MSE_f} \quad (5)$$

where MSE_u denotes the mean squared error loss corresponding to the initial data, N_u is the total number of training data, MSE_f is the mean squared error at a finite set of collocation points and N_f is the total number of collocation points. The

number of collocation points and training data influence the prediction accuracy and the computational time to optimize the loss function. The error MSE_u enforces the boundary conditions of the independent variables x and MSE_f enforces the physics of the dynamical system imposed by the condition (3), i.e. it penalizes deviations of the predicted physical law. Given a training data set and known system parameters λ, we seek to find the parameters (weights and biases) of the neural networks which minimize (5). If the parameters λ are unknown, we train for the same objective but consider the system parameters as additional variables.

1) Physics-informed neural networks capturing power system dynamics: We show how physics-informed neural networks can be used to derive δ and $\omega = \dot{\delta}$ of the swing equation (2) at any time instant t and for a range of mechanical power P_1. We assume that the system parameters $\lambda := \{m_1, d_1, B_{12}\}$ are known and the voltages V_1 and V_2 are fixed. As a result the system input is defined as $x := \{P_1\}$. In contrast with conventional numerical solvers, which require the conversion of higher-order ordinary differential equations (ODEs) to first-order in order to solve them (by introducing additional variables), physics-informed neural networks can directly incorporate higher-order ODEs, as we show in (7).

Incorporating (2) to the neural network, function (4) is given by:

$$u(t, x) := \delta(t, P_1), \tag{6}$$

$$f_\delta(t, P_1) = m_1\ddot{\delta} + d_1\dot{\delta} + B_{12}V_1V_2\sin(\delta) - P_1, \quad P_1 \in [P_{\min}, P_{\max}], \quad t \in [0, T] \tag{7}$$

The interval $[0, T]$ can be defined based on the time period of interest for the dynamic simulation. The domain Ω of the input P_1 is restricted to $[P_{\min}, P_{\max}]$. The neural network output is $\delta(t, P_1)$. After the training phase, the frequency signal $\omega := \dot{\delta}$ is extracted as a function of the estimated angle δ. As a result, the prediction error of the frequency ω depends on the prediction error of the angle δ and the differential method. In the rest of the paper, we refer to this neural network structure as NN_δ.

2) Data-driven discovery of inertia and damping coefficients: Information about power system parameters such as system inertia is of significant importance for system operators to prevent large frequency deviations and maintain frequency stability. As described in [16], due to varying generation of converter-connected renewable energy sources, the inertia level of power systems becomes uncertain and has to be estimated (or predicted) at regular time intervals [17]. Physics-informed neural networks can be used to address the problem of system identification and data-driven discovery of partial differential equations. For this case, we define m_1 and d_1 as unknown parameters in (7). The structure of the physics-informed neural network remains the same, with the only difference that a subset of the system parameters λ are now treated as additional variables when minimizing (5) during neural network training.

III. Simulation & Results

A. Simulation Setup

Besides an initial training set, to assess the neural network performance we also need an extensive test data set. To create the training and test data sets we use the numerical solver $ode45$ in MATLAB with a time step of $0.1s$ and time interval $T = [0, 20s]$, resulting in 201 time steps for each trajectory. The voltage magnitudes V_1 and V_2 are equal to 1 p.u. and B_{12} = 0.2 p.u. In our first case study, we assume system inertia and damping are known, and that the system is not at an equilibrium. Assuming an uncertain active power input in the range $P_1 = [0.08, 0.18]$ and initial values for δ and ω equal to 0.1 rad and 0.1 rad/s, we generate 100 trajectories. As a result, our entire test and training dataset consists of $20'100$ samples. We consider the interval from $[0.08, 0.18]$ to show the capability of the physics-informed neural network to accurately predict trajectories for uncertain power injections. For values larger than 0.18, the system becomes unstable, and for values lower than 0.08 multiple oscillations occur. For these regimes, we observe lower prediction accuracy, and different trained physics-informed neural networks could be used to achieve high accuracy in each of these regimes.

In our second case study, inertia and damping are also unknown parameters. Given scattered observed data about active power, frequency and angle measurements, our goal is to identify the parameters m_1 and d_1 of (7), as well as to obtain the trajectory of δ. Considering that the levels of inertia and damping vary, we assign 10 different values to m_1 and d_1 that lie within the range of $[0.1, 0.4]$ and $[0.05, 0.15]$, respectively. To this end, for each of the 10 pairs $\{m_1, d_1\}$ we generate 40 trajectories.

Next, before starting the training procedure, as usual for neural networks, we need to determine an appropriate number of hidden layers and number of neurons per layer, the amount of training data N_u and the number of collocation points N_f. We carried an extensive investigation of the appropriate values for each of those parameters, assessing the relative L_2 error between the predicted and the exact solution of $\delta(t, P_1)$ and $\omega(t, P_1)$ for a range of different configurations. In the case studies, we report results only for the most suitable configuration, which achieved the lowest L_2 error. Similar to [10], as the required amount of training data N_u is very small (only 40 data points), we use a gradient-based optimization algorithm to optimize the loss function $MSE = MSE_u + MSE_f$ in (5). We perform neural network training and testing in TensorFlow on a laptop (Intel Core i7 3.9 GHz, 32-GB RAM, single NVIDIA GeForce 940MX 2-GB). The hidden layers of the neural network use hyperbolic tangent activation functions. The code to reproduce the results is available online [14].

B. Data-driven solution of frequency dynamics through physics-informed neural networks

The following parameters were selected to obtain the lowest L_2 error on the test data: we select a set of N_u = 40 randomly distributed initial and boundary data across the entire spatio-temporal domain, N_f = $8'000$ collocation points, and a 5-layer neural network with 10 neurons per hidden layer. Observe that compared to conventional neural network approaches, we only need a very small amount of samples (N_u = 40). Increasing N_u in our simulations, led to over-fitting to the training data. Training took 223 seconds and the relative L_2 error between exact and predicted solutions on the $11'600$-points test dataset is $1.34 \cdot 10^{-2}$. Fig. 3 depicts the comparison between the predicted and the actual trajectory of the angle $\delta(t)$ and the frequency $\omega(t)$. The best and worst $\{\delta, \omega\}$ estimation during different active power inputs P_1 in terms of L_2 error on both

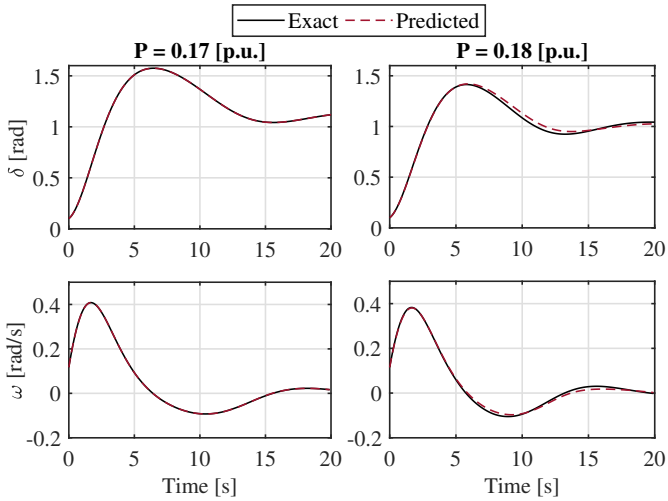

Fig. 3. Comparison of the predicted and exact solution for the angle $\delta(t)$ and frequency $\omega(t)$ with the physics-informed neural network NN_δ. Note that to compute the frequency $\omega(t)$ we perform numerical differentiation of the angle $\delta(t)$ using a Newton method. In the left figures, we show the most accurate estimation of the trajectory of $\delta(t)$ and $\omega(t)$, with a relative L_2 error of $2.37 \cdot 10^{-2}$. In the right figures, we show the least accurate estimation of the trajectory of $\delta(t)$ and $\omega(t)$, with a relative L_2 error of $2.55 \cdot 10^{-4}$.

training and test sets are depicted in the left and right side of the figure, respectively. To extract the frequency ω we differentiate the signal associated with the angle δ. To this end, we numerically approximate the derivative of a function as: $\omega(t) = \lim_{h\to 0} \frac{\delta(t+h)-\delta(t)}{h}$. The value of h depends on the simulation time step. In this study, we generated the trajectories with a fixed step of $h = 0.1s$. In future work, we will use automatic differentiation [11] to extract the frequency directly from the physics-informed neural network. It can be observed that the physics-informed neural network is able to predict the trajectory of the angle $\delta(t)$ with high accuracy, and that the frequency signal $\omega(t)$ can be successfully recovered using numerical differentiation.

After training, we evaluate the neural network performance in terms of computational speed required for solving the differential equation defined by (2). For 100 different initial conditions of (2), the $ode45$ solver takes on average 0.45 s to solve the differential equations and the neural network only 0.016 s, resulting to a speed-up of factor of 28. We expect that for larger systems the computational speed-up will be even higher, as solving large-scale differential equations is computationally very expensive, whereas the evaluation of a trained neural network remains computationally low even for large network sizes. Additionally, and most importantly, the physics-informed neural network can directly determine δ at any specified time step $\delta(t_1, P_1)$, whereas numerical methods always have to start integrating from the boundary conditions at $t = t_0$ until they reach $t = t_1$. The computational time for evaluating any random time step (e.g. at $t_1 = 10$ s) is $4 \cdot 10^{-3}$ s, whereas integrating from $t_0 = 0s$ up to $t_1 = 10$ s with the $ode45$ solver takes 0.35 s, resulting to a *substantial speed-up of almost two order of magnitude* for the physics-informed neural network (87 times to be exact). This illustrates the capability of physics-informed neural networks to predict directly the solution to higher-order differential equations with high accuracy and low computational cost, offering significant advantages over classical numerical integration tools.

1) Predicting both angle δ and frequency ω as separate neural network outputs: Within our investigations, we also attempted to train a physics-informed neural network that considers δ and ω as separate outputs, essentially setting $f_\omega = \dot{\delta} - \omega$ and $f_\delta = m_1\dot{\omega} + d_1\omega + B_{12}V_1V_2\sin(\delta) - P_1$. To obtain the lowest L_2 error in this case, we had to select again a set of N_u = 40 randomly distributed initial and boundary data, a 5-layer neural network with 10 neurons per hidden layer, but a set N_f = $50'000$ collocation point (instead of $8'000$ in the previous case) . The model training took approximately 30 minutes as more collocation points N_f are required to obtain a satisfactory prediction error. Considering that $\delta(t)$ and $\omega(t)$ are predicted as separate outputs, the relative L_2 errors between the exact and predicted solutions are $9.43{\cdot}10^{-2}$ and $1.51{\cdot}10^{-1}$, respectively, and are higher than for the NN_δ structure. It becomes obvious that the neural network architecture with the single output δ (and subsequent numerical differentiation to determine ω) is preferable in terms of training time and predictive accuracy.

C. Data-driven discovery of inertia and damping coefficients through physics-informed neural networks

In this subsection, we evaluate the performance of the physics-informed neural network to predict system inertia and damping from observed trajectories. In this case study, we assume that m_1 and d_1 are unknown, and instead we have a set of limited training datapoints $\{t, P_1, \delta\}$. Contrary to the usual practice of first training a neural network and then using it, our objective here is exploit the physics-informed neural network training procedure to determine m_1 and d_1. To illustrate the effectiveness of this approach, we perform this analysis for 10 different pairs of $\{m_1, d_1\}$ and evaluate the average predictive accuracy. We select a set of N_u = 100 randomly distributed points across the spatio-temporal domain from the exact solutions of (2) for each inertia level. A 5-layer neural network with 30 neurons per hidden layer is trained for each inertia level with the corresponding trajectories in order to predict the system parameters and $\delta(t)$. The resulting average errors for predicting m_1 and d_1 over the 10 different cases are 0.74% and 1.28%, respectively. The average training time of the neural network to identify the system parameters was less than 60 seconds. This means that with a limited training dataset, and within 60 seconds, we can accurately predict the inertia and damping level of a system. Considering that the swing equation (2) is often used to approximate the aggregate dynamic behavior of large power systems, these results demonstrate that physics-informed neural network show substantial potential to not only accurately derive δ and ω but also predict both system inertia and damping. Last but not least, the relative L_2 errors between the exact and predicted solutions for the phase angle are less than 10^{-1} over the 10 different cases of $\{m_1, d_1\}$. This shows the potential of physics informed neural networks to be used as a dynamic state estimator, when the model parameters are unknown [13].

IV. Discussion and Outlook

This work introduces for the first time in power systems a neural network training procedure that explicitly considers the underlying differential and algebraic equations describing power system behavior. This unlocks a series of opportunities in power systems, as physics-informed neural networks may

be able to accurately determine the solution of differential-algebraic sets of equations several orders of magnitude faster than traditional methods relying on numerical integration. Still, to unlock this potential, there are several challenges to be addressed.

a) Number of training data: Besides the limited number of training data, physics-informed neural networks as described in this paper need to generate a substantial number of collocation points. In our case studies, we used $N_u = 40$ points as input data and $N_f = 8'000$ collocation points. It is expected that for larger systems, a much larger number of collocation points will be necessary, which will result to a longer training time. In our future work, we plan to investigate methods using Runge-Kutte integration schemes such as the ones proposed in [10] which can eliminate the need for collocation points.

b) Scalability: Although the swing equation is a good first approximation for first-swing instability, and single-machine infinite-bus systems are still used as aggregate models of large power systems, we still need to explore what are the computational needs if we were to apply these methods in large scale power systems and how to address the associated challenges related to the neural network training. Particularly, the comparison with numerical solvers and approximation techniques like polynomial fits will serve as a benchmark.

c) Range of applications: As shown in this paper, physics-informed neural networks can determine two orders of magnitude faster the rotor angle and frequency at any time instant for uncertain power inputs. At the same time, they can accurately identify uncertain parameters such as inertia and damping. Future applications must also assess cases that include both stable and unstable equilibria, a wide range of different dynamic phenomena, including small-signal stability, voltage stability and converter dynamics [18], discrete events, such as protection actions, as well as power system optimization, among numerous others. In our simulation study, we observed high accuracy for a single stable swing prediction, but for different regimes such as multiple oscillations or unstable conditions, different physics-informed neural networks might have to be trained. We also need to examine if such neural networks can capture discrete events, such as protection actions, or if we need to develop a hybrid approach, using physics-informed neural networks as a numerical solver only during the continuous dynamics before and after a discrete event. For power system applications, physics-informed neural networks can (and should) be combined with neural network verification methods, see [19]. In this way, they would no longer be considered a black box, but instead we would be able to extract formal guarantees for their behavior.

V. Conclusions

To the best of our knowledge, for power system applications, this is the first paper to propose physics-informed neural networks. Explicitly considering the power system governing equations, we are able to determine the solution of differential-algebraic systems of equations at a fraction of the time required for conventional numerical approaches. Physics-informed neural networks require substantially less training data, while achieving high accuracy, due to the inclusion of the underlying swing equation. This paper introduces the general framework and presents results for a single-machine infinite-bus system. In our case studies, we demonstrate how physics-informed neural networks can accurately determine the rotor angle and frequency *87 times* faster than conventional numerical methods. We further demonstrate their successful identification of uncertain system parameters such as inertia and damping from a limited set of input data. Our results showcase the potential for successful application of these methods in larger systems, unlocking a series of opportunities for power system security and optimization, achieving good accuracy and high computational speed. Future work will explore a series of possible applications and potential improvements in the training procedure.

Acknowledgement

This work is supported by the multiDC project funded by Innovation Fund Denmark, Grant No. 6154-00020B.

References

[1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[2] L. A. Wehenkel, *Automatic learning techniques in power systems*. Springer Science & Business Media, 2012.

[3] B. Donnot, I. Guyon, M. Schoenauer, P. Panciatici, and A. Marot, "Introducing machine learning for power system operation support," *CoRR, arXiv preprint arXiv:1709.09527*, 2017.

[4] M. Sun, I. Konstantelos, and G. Strbac, "A deep learning-based feature extraction framework for system security assessment," *IEEE Transactions on Smart Grid*, vol. 10, no. 5, pp. 5007–5020, Sep. 2019.

[5] J. H. Arteaga, F. Hancharou, F. Thams, and S. Chatzivasileiadis, "Deep learning for power system security assessment," in *2019 IEEE Milan PowerTech*, June 2019, pp. 1–6.

[6] F. Fioretto, T. W. K. Mak, and P. V. Hentenryck, "Predicting ac optimal power flows: Combining deep learning and lagrangian dual methods," *arXiv preprint arXiv:1909.10461*, 2019.

[7] F. Thams, A. Venzke, R. Eriksson, and S. Chatzivasileiadis, "Efficient database generation for data-driven security assessment of power systems," *IEEE Transactions on Power Systems*, pp. 1–1, 2019.

[8] A. Venzke, D. K. Molzahn, and S. Chatzivasileiadis, "Efficient creation of datasets for data-driven power system applications," *arXiv preprint arXiv:1910.01794*, 2019.

[9] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Advances in neural information processing systems*, 2018, pp. 6571–6583.

[10] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686 – 707, 2019.

[11] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic differentiation in machine learning: a survey," *Journal of machine learning research*, vol. 18, no. 153, 2018.

[12] M. Abadi *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.

[13] J. Zhao *et al.*, "Power system dynamic state estimation: Motivations, definitions, methodologies, and future work," *IEEE Trans. on Power Syst.*, vol. 34, no. 4, July 2019.

[14] G. S. Misyris, A. Venzke, and S. Chatzivasileiadis, "Online Appendix: Physics-Informed Neural Networks for Power Systems," 2019. [Online]. Available: https://github.com/gmisy/Physics-Informed-Neural-Networks-for-Power-Systems/

[15] S. Chatzivasileiadis, T. L. Vu, and K. Turitsyn, "Remedial actions to enhance stability of low-inertia systems," in *IEEE Power and Energy Society General Meeting 2016, Boston, MA, USA*, July 2016, pp. 1 –5.

[16] G. S. Misyris, S. Chatzivasileiadis, and T. Weckesser, "Robust frequency control for varying inertia power systems," in *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe*, 2018, pp. 1–6.

[17] D. Zografos and M. Ghandhari, "Power system inertia estimation by approaching load power change after a disturbance," in *2017 IEEE Power and Energy Society General Meeting*, 2017, pp. 1–5.

[18] G. S. Misyris, J. A. Mermet-Guyennet, S. Chatzivasileiadis, and T. Weckesser, "Grid supporting vscs in power systems with varying inertia and short-circuit capacity," in *2019 IEEE Milan PowerTech*, June 2019, pp. 1–6.

[19] A. Venzke and S. Chatzivasileiadis, "Verification of neural network behaviour: Formal guarantees for power system applications," *arXiv preprint arXiv:1910.01624*, 2019.