Surrogate Model Training Data for FIDVR-related Voltage Control in Large-scale Power Grids

Tianzhixi Yin
Pacific Northwest National Laboratory
Richland, WA 99354
tianzhixi.yin@pnnl.gov

Ramij R. Hossain
Pacific Northwest National Laboratory
Richland, WA 99354
ramijraja.hossain@pnnl.gov

Qiuhua Huang*
Pacific Northwest National Laboratory
Richland, WA 99354
qiuhuahuang@mines.edu

Renke Huang*
Pacific Northwest National Laboratory
Richland, WA 99354
huangrenke@gmail.com

Jie Tan
Google
Mountain View, CA 94043
jietan@google.com

Wenhao Yu
Google
Mountain View, CA 94043
magicmelon@google.com

Abstract

This work presents an effective machine learning (ML) data set related to the short-term voltage dynamics in power systems. Power systems dynamics are highly nonlinear and intricate. Model designs/specifications in power systems need expertise to capture dynamic phenomena. ML has become an important tool for analyzing complex behaviors of physical systems, but ML models need quality data sets for training and testing. Learning surrogate models to replicate certain dynamic behaviors of power systems is a growing area of interest; however, building required data sets can be challenging. We utilize the high performance computing (HPC)-based grid simulator GridPACK to create voltage dynamics of a bulk power system, namely the IEEE 300 bus test system, capturing fault-induced delayed voltage recovery (FIDVR) phenomenon. This FIDVR is generally mitigated by the under voltage load shedding (UVLS)-based control strategy. The data set created here contains the trajectory data of voltage dynamics under different control actions generated by standard UVLS strategy and random noise. We present the structure of the data set and its application in learning a dynamic surrogate model. Finally, other suitable ML-based applications of the given data set are discussed, thereby helping to strengthen reusable science practices. The dataset is open-sourced at `https://github.com/pnnl/MBDRL/tree/master/surrogate_model_data`.

1 Introduction

In recent years, power and energy sector has seen steep rise of machine learning methods [Entezari et al., 2023] in various application sub-fields ranging from load forecasting, state estimation, anomaly detection to real-time decision making [Forootan et al., 2022, Xie et al., 2020]. The data set presented

*Qiuhua Huang and Renke Huang were at PNNL when conducting this research.

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

in this paper focuses on the dynamical side of power systems, which follows a set of differential-algebraic equations (DAE). With system size, e.g., the number of buses/nodes, generators, and other components, the dimension and complexity of the model increase manifold, making model-based optimization and decision-making computationally challenging; whereas machine learning methods have the capability of learning from data with minimal dependence on mathematical models, thereby providing an effective solution to tackle the issues of high-dimensional models of physical systems and associated computational challenges.

The success of machine learning methods primarily depends on good-quality data. The availability of large data set related dynamic events from real-world power grid are constrained due to many factors, e.g. safety-critical nature, low occurrence, limitations of experiments. Fault-induced delayed voltage recovery (FIDVR) is a dynamic phenomena observed after fault clearance due to stalled residential air conditioning (AC) units (powered by single-phase induction motors), and causes delayed voltage recovery, which may lead to widespread cascading blackouts Eto [2015]. This short-term voltage instability can be mitigated by controlled shedding of loads achieved through under voltage load shedding (UVLS) relays Lu et al. [2009]. UVLS is rule-based method and may fail to adapt with the operational changes. To improve the solution aspects, voltage recovery dynamics with control actions needs to be studied. But due to scarcity of the real-world data, the use of physics-based simulators is a common practice in the power systems research community. In general, FIDVR simulation for the IEEE 300 bus-systems needs modeling expertise. Even with GridPACK, a HPC-based simulator Palmer et al. [2014], FIDVR simulation can take several seconds, and is not compatible with real-time fast decision-making. A data-trained surrogate model replicating the voltage recovery dynamics can break this barrier, as the inference time of machine learning models is usually much less compared to physics-based simulation. In this paper, we simulate the FIDVR phenomenon with bulk IEEE 300-bus systems, under standard UVLS control blended with random noise. The simulated voltage trajectory data is then converted to data tuple considering multi-step transition data format (discussed later), and utilized to learn a surrogate model. Further, we present a surrogate model training example to show the application of the data set.

2 Power System Dynamic Simulation, Test Systems, and Computational Challenges

Power system dynamics subjective to external disturbances and controls can be formulated as a set of differential algebraic equations (DAEs) as follows:

$$\dot{\mathbf{x}}_t = f(\mathbf{x}_t, \mathbf{y}_t, d_t, \mathbf{a}_t,), \mathbf{x}_{t_0} = \mathbf{x_0}, \tag{1a}$$

$$0 = g(\mathbf{x}_t, \mathbf{y}_t, d_t, \mathbf{a}_t), \mathbf{y}_{t_0} = \mathbf{y_0}, \tag{1b}$$

where 1a represents system dynamics; $\mathbf{x_t}$ represents dynamic state variables, such as the generator rotor angles and speeds, etc.; 1b represents power balance at each bus; $\mathbf{y_t}$ represents the algebraic state variables of the power grids such as voltages at nodes (or buses) of the grid; $\mathbf{x_0}$ and $\mathbf{y_0}$ are initial values;$\mathbf{a_t}$ represents control actions such as load shedding; and d_t represents the disturbance.

For large-scale power systems, the dimension of the DAEs is high and the problem is computationally intensive to solve. The GridPACK Palmer et al. [2014], developed by the Pacific Northwest National Laboratory (PNNL), is a notable open-source software platform tailored for power grid system simulation and analysis. One of the prominent features of GridPACK is its inherent capability to harness the computational muscle of HPC resources efficiently.

Even with HPC resources, it is still challenging to run a large number of power system dynamic simulations for training reinforcement learning (RL) agents for controlling power systemsHuang et al. [2022]. This computational demand is primarily attributed to the exhaustive interaction with the environment required by most RL algorithms. Given these computational constraints, there is an emerging imperative to engineer a surrogate model, which can emulate power system dynamics, thereby optimizing the computational efficiency.

The IEEE 300-bus system is a widely recognized benchmark power system network in the field of electrical engineering and power systems research. It serves as a valuable tool for evaluating various power system analysis and optimization techniques. This system features 300 buses or substations interconnected by transmission lines and transformers.

3 The Data Set

In our approach, we employed a predetermined rule-based policy, specifically the UVLS protocol, augmented with stochastic noise, to synthesize system trajectories across diverse operational scenarios for training. The resultant offline dataset was constructed through preprocessing of trajectory data, organized into multi-step transition tuples. UVLS is a systematic strategy designed to maintain grid stability. When the system experiences undervoltage conditions, which can jeopardize the integrity and reliability of the power system, UVLS is activated to deliberately disconnect specific loads. This action is predicated on a hierarchy of predetermined setpoints, ensuring that only non-essential loads are shed first, thereby preventing widespread blackouts and mitigating potential damages to the grid infrastructure.

We curated the offline dataset by leveraging the following methodologies:

(1) Employing an UVLS policy and subsequently infusing stochastic perturbations, in the form of random Gaussian noise (zero-mean with a standard deviation of 0.1), to the resultant actions (from -0.2 to 0, representing percentages of load shed).

(2) Integrating power flow and fault scenarios, specifically designed for the training duration.

(3) Employing a GridPACK-facilitated time-domain simulation to yield voltage trajectory data, governed by the action sequences from the UVLS policy augmented with random noise.

(4) Formulating the M-step transition tuple, $\{s_t, a_t, s_{t+1}, a_{t+1}, \cdots, s_{t+M}\}$, derived from the aforementioned voltage trajectory data, and subsequently amalgamating it into the offline dataset, $\mathcal{D}_{\text{offline}}$.

The dataset comprises inputs (X) and corresponding outputs (y), with X being structured as a multidimensional array of dimensions $600 \times 80 \times 180$. The first dimension (600) represents the total number of samples in our dataset, each corresponding to a different combination of fault location and random noise. We have considered six distinct fault locations in our study. The second dimension (80) represents the number of discrete time steps observed during each simulation, referred to as "rollouts". These time steps capture the dynamic behavior of the power system. The third dimension (180) of X is further divided into various components, including:

(1) **108-bus Voltage Measurements**: This subcomponent encompasses voltage measurements taken across the power system at different buses. Figure 1 shows an example of voltages from ten buses during fault 1.

(2) **34 Load Bus Actions**: Within each time step, actions are specified for 34 load buses. These actions pertain to the manipulation or control of loads at these buses.

(3) **Additional Information**: For each time step, four supplementary pieces of information are provided: (i) Start time of the fault; (ii)Duration of the fault; (iii) Step size of the fault; (iv) The time distance from the current step to the start of the fault.

Figure 1: Example of Voltages during Fault 1.

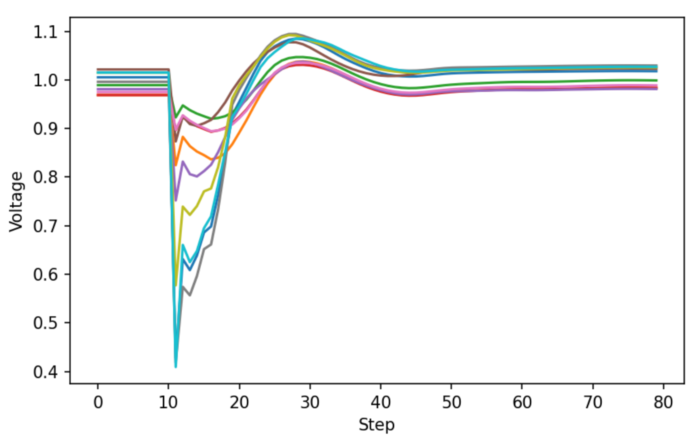

As for the outputs (y), they consist of voltage measurements across the power system, specifically encompassing 108 voltage values. These voltage measurements reflect the electrical state of the power grid at each corresponding time step.

In summary, our dataset is structured to encompass diverse fault scenarios, each with specific fault locations and random noise variations. The input data (X) encapsulates dynamic information about the power system, including voltage measurements, load bus actions, and auxiliary fault-related details. The output data (y) comprises voltage measurements that serve as indicators of the system's response to these fault scenarios and dynamic conditions. This dataset serves as a valuable resource for the analysis and modeling of power system behavior under fault conditions, offering insights into the system's dynamic responses.

4 ML-based Application

In our research, we harness the capabilities of deep neural networks (DNNs) to model the transition dynamics in power system simulations. Specifically, we employ a DNN denoted as $f_\phi(s_t, a_t)$. Our DNN architecture comprises three layers with respective hidden units of 1000, 500, and 200. Rectified Linear Units (ReLU) serve as the activation function throughout the layers. The activation function for the output layer is Sigmoid. Conventionally, the next state s_{t+1} is directly inferred from the current state s_t and the taken action a_t. Thus, we represent this relation as $s_{t+1} = f_\phi(s_t, a_t)$. Given a compilation of state-transition tuples, denoted as $\mathcal{D} = \{(s_t, a_t, s_{t+1})\}$, the training of $f_\phi(\cdot, \cdot)$ is done by minimizing the loss function:

$$\mathcal{L}(\phi) = \frac{1}{|\mathcal{D}|} \sum_{(s_t, a_t, s_{t+1}) \in \mathcal{D}} \| s_{t+1} - f_\phi(s_t, a_t) \|_2^2 \tag{2}$$

The above mentioned model formulation, given by $s_{t+1} = f_\phi(s_t, a_t)$, exhibits computational inefficiencies when s_t and s_{t+1} are proximal in the state-space, as elucidated in Nagabandi et al. [2018]. Consequently, this model configuration, $f_\phi(s_t, a_t)$, tends to exhibit inaccuracies in predicting future states. This challenge becomes more pronounced under: (a) intricate dynamics, exemplified in power systems, and (b) minute temporal intervals between sequential time steps. Furthermore, within the context of Deep Reinforcement Learning (DRL), it becomes imperative to accurately propagate dynamics across extended horizon rollouts. Any discrepancies in these dynamics culminate in the amplification of error throughout the said horizon. In response to these challenges, our approach incorporates: (a) the differential between successive states while defining $f_\phi(\cdot, \cdot)$, expressed as $s_{t+1} - s_t = f_\phi(s_t, a_t)$, and (b) a multi-step loss function, represented by (3). This stands in contrast to the traditional single-step loss function, defined in (2). For a structured formulation of the M-step loss-driven training of $f_\phi(\cdot, \cdot)$, we adhere to the methodology delineated in Yang et al. [2020]. We follow the steps below to formulate the multi-step (M-step) loss based training of $f_\phi(\cdot, \cdot)$:

(a) Construct the M-step ground-truth transition tuple, denoted as $\mathcal{D}_M = \{(s_t, a_t, s_{t+1}, a_{t+1}, \cdots, s_{t+M})\}$, which encapsulates state and action sequences over the designated horizon.

(b) Forecast the future state sequence commencing from s_t, by employing the recurrence relation:

$$\hat{s}_{t+\tau+1} = \hat{s}_{t+\tau} + f_\phi(\hat{s}_{t+\tau}, a_{t+\tau})$$

spanning from $\tau = 0$ to M. It is imperative to emphasize that, barring the inaugural iteration, the DNN $f_\phi(\cdot, \cdot)$ utilizes the estimated states $\hat{s}_{t+\tau}$ as opposed to the predicted states $s_{t+\tau}$ for its inputs. Incorporating estimated states as successive inputs to the DNN aids in attenuating the amplification of prediction errors across extended horizons.

(c) The DNN undergoes a training regimen employing mini-batch stochastic gradient descent, with the objective of minimizing the loss as articulated in (3).

$$\mathcal{L}_{ms}(\phi) = \frac{1}{|\mathcal{D}_M| \times M} \sum_{\substack{(s_{t:t+M}, \\ a_{t:t+M-1}) \\ \in \mathcal{D}_M}} \sum_{\tau=0}^{M-1} \left\| s_{t+\tau+1} - \hat{s}_{t+\tau+1} \right\|_2^2 \tag{3}$$

In the experimental setup, we selected a 5-step loss configuration, denoted as $M = 5$. The DNN undergoes an optimization phase employing the mini-batch stochastic gradient descent algorithm, facilitated by the ADAM optimizer. The hyperparameters for this phase are as follows: learning rate $\alpha = 0.001$, first moment coefficient $\beta_1 = 0.9$, second moment coefficient $\beta_2 = 0.999$, and smoothing term $\epsilon = 1 \times 10^{-7}$. The surrogate model performance on the offline data set is shown in Table 1. We randomly selected 10% of the fault cases to be the testing set and 9% of the cases to be the validation set.

Table 1: Mean Square Errors for the Surrogate Model across Different Phases

Phase	**MSE**
Training	2.0252×10^{-7}
Validation	1.0594×10^{-7}
Testing	2.6207×10^{-7}

Comparative analysis indicates that the surrogate model executes approximately an order of magnitude faster than its GridPACK counterpart. It is salient to highlight that this computational disparity is anticipated to widen commensurately with the upscaling of power system dimensions.

5 Prospective Applications

The use of surrogate models in our approach significantly enhances the efficiency of model-based reinforcement learning training. Traditional methods often suffer from being both time-consuming and requiring extensive sampling, which can be infeasible, especially in complex scenarios like power grid management. In contrast, our surrogate model-driven approach allows for more rapid insights and actionable outcomes with fewer data samples. This efficiency has profound implications for various applications. For instance, in the realm of grid stabilization and reliability, it enables real-time decisions that help stabilize the grid during unpredictable fluctuations. Similarly, as the integration of fluctuating renewable energy sources becomes increasingly common, our method offers a viable solution to predict and manage these variations, ensuring a fine balance between renewable inputs and grid stability.

6 Conclusion

In this work, we have presented a valuable dataset stemming from power system simulators, specifically curated for the advancement of surrogate model-driven reinforcement learning in power grid voltage control. We believe this dataset will pave the way for innovative solutions in grid management, addressing both current challenges and future complexities. We encourage researchers and practitioners to leverage this resource, fostering collaborative advancements in the realm of sustainable and efficient power systems.

References

A. Entezari, A. Aslani, R. Zahedi, and Y. Noorollahi. Artificial intelligence and machine learning in energy systems: A bibliographic perspective. *Energy Strategy Reviews*, 45:101017, 2023.

J. Eto. Area of interest 1: Interregional electricity reliability issue and assessment analysis. 2015.

M. M. Forootan, I. Larki, R. Zahedi, and A. Ahmadi. Machine learning and deep learning in energy systems: A review. *Sustainability*, 14(8):4832, 2022.

R. Huang, Y. Chen, T. Yin, X. Li, A. Li, J. Tan, W. Yu, Y. Liu, and Q. Huang. Accelerated derivative-free deep reinforcement learning for large-scale grid emergency voltage control. *IEEE Transactions on Power Systems*, 37(1):14–25, 2022. doi:10.1109/TPWRS.2021.3095179.

N. Lu, B. Yang, Z. Huang, and R. Bravo. The system impact of air conditioner under-voltage protection schemes. In *2009 IEEE/PES Power Systems Conference and Exposition*, pages 1–8, 2009. doi:10.1109/PSCE.2009.4840053.

A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7559–7566, 2018. doi:10.1109/ICRA.2018.8463189.

B. Palmer, W. Perkins, Y. Chen, S. Jin, D. Callahan, K. Glass, R. Diao, M. Rice, S. Elbert, M. Vallem, and Z. Huang. Gridpack: A framework for developing power grid simulations on high performance computing platforms. In *2014 Fourth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing*, pages 68–77, 2014. doi:10.1109/WOLFHPC.2014.12.

J. Xie, I. Alvarez-Fernandez, and W. Sun. A review of machine learning applications in power system resilience. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2020.

Y. Yang, K. Caluwaerts, A. Iscen, T. Zhang, J. Tan, and V. Sindhwani. Data efficient reinforcement learning for legged robots. In L. P. Kaelbling, D. Kragic, and K. Sugiura, editors, *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 1–10. PMLR, 30 Oct–01 Nov 2020. URL `https://proceedings.mlr.press/v100/yang20a.html`.