Digital Twin Generalization with Meta and Geometric Deep Learning

Raffael Theiler, Olga Fink
ENAC IIC IMOS
École Polytechnique Fédérale de Lausanne (EPFL)
Lausanne, CH-1015, Lausanne
{raffael.theiler, olga.fink}@epfl.ch

ABSTRACT

Deep digital twins (DDTs) are deep neural networks that encode the behavior of complex physical systems. DDTs are excellent system representations due to their ability to continuously adapt to operational changes and their capability to capture complex relationships between system components and processes that cannot be explicitly modeled. For this challenge, DDTs benefit greatly from recent success in geometric deep learning (GDL) which allows the integration of information from multiple systems based on schematic representations. A major challenge in training DDTs is their dependence on the quality and representativeness of training data, especially under the dynamic conditions typical in prognostics and health management (PHM). Recent developments in differentiable simulation present new opportunities for optimizing the training data representativeness. In this thesis, we propose a novel meta-learning framework that trains DDTs using the output from differentiable simulators. This setup enables active optimization of training data sampling through gradient computation, enhancing training speed, robustness, and data representativeness. We extend this framework to address challenges in multi-system data integration in power grids and fault detection in railway traction networks. By applying our framework, we aim to tackle significant challenges in forecasting, anomaly detection and sensor-fault analysis using advanced data fusion techniques. Our approach promises substantial improvements in DDT robustness and operational efficiency, with its effectiveness to be demonstrated through empirical studies on both simple and complex case studies within the power systems domain.

1. INTRODUCTION

Power systems are a paramount example of complex systems (Cuadra, Salcedo-Sanz, Del Ser, Jiménez-Fernández, & Geem, 2015). Conventional and data-driven methods have been developed for the analysis of power systems (Liao, Bak-Jensen, Radhakrishna Pillai, Wang, & Wang, 2022), and other industrial assets (Hu, Miao, Si, Pan, & Zio, 2022) over many years. Yet, they suffer from non-stationary behaviors of complex systems such as changes in environmental and operational conditions (Soleimani, Campean, & Neagu, 2021). Such non-stationarity poses a challenge to traditional data-driven machine learning which is trained on the i.i.d assumption that data are identically and independently distributed with balanced training and testing sets (Jardine, Lin, & Banjevic, 2006). However, the requirement for balanced amounts of healthy and unhealthy data across different operational states cannot be met in the industrial environment due to the long operational lifespans of critical assets and the challenges in storing and processing the associated data (Booyse, Wilke, & Heyns, 2020). The resulting imbalance violates the i.i.d assumption. Therefore, sophisticated approaches for the operation and maintenance of industrial systems with data-driven models are required to ensure robust and reliable operation. Enhancing the generalization ability of the data-driven model with respect to these conditions becomes essential in both industry and academic fields (Wang et al., 2022).

For operation and maintenance, Digital twins (DT) have been one of the emerging tools applied for industrial systems. They are virtual representations of these industrial systems that can be used to approximate the behavior in a digital environment (Booyse et al., 2020). DTs are different from data-driven simulators because they make use of additional (real-time) sensory information from the physical system for online updates of the system state to real operational and environmental scenarios. The most prominent application of DT is for control. However, they have been increasingly applied for fault detection, diagnostics, and prognostics where DTs continuously represent the system to help engineers better identify deviations from expected behavior to take corrective action. Furthermore, DTs are used for prognostics and health management (PHM), where the DT can approximate the evolution of system conditions, predict when maintenance is needed, and plan it accordingly (Aivaliotis, Georgoulias, Arkouli, & Makris, 2019). Advances in machine learning (ML) and specifically deep learning (DL), have led to the development of neural DTs and neural simulators (Liao et al., 2022). We refer to those neural-network-driven DTs as deep digital twins (DDT) (Booyse et al., 2020).

The ability to train models of complex industrial systems on real-world data enables new learning strategies. Geometric deep learning (GDL) can be leveraged to learn currently hidden implicit relations between neighboring sub-systems on a larger scale. System connectivity that cannot be explicitly modeled with current tools can be determined. and information aggregation into common latent representations can be advanced. The training data does not have to be limited to recorded real-world data, but can also be generated synthetically in simulation, which drastically reduces the total cost of training in certain safety-critical environments, e.g. in robotics.

In this thesis, we pursue two avenues to improve the robustness and generalization of DDT. First, recent literature suggests that synthetic data is a powerful and currently underdeveloped method to improve the generalization ability of DL models (Zhou, Liu, Qiao, Xiang, & Loy, 2022), its potential has not yet been exploited to improve the robustness and generalization of DDTs. We envision one of the potential avenues can be to train a DDT on data coming from a preceding data-driven simulator. Furthermore, a model that is trained on synthetic data can later be efficiently fine-tuned to further improve the performance (Biggio, Bendinelli, Kulkarni, & Fink, 2022). In this novel two-stage approach, comprising a data-driven simulator and a subsequent DDT, the generalization ability of the DDT can be significantly improved by sampling synthetic data. This concept paves an avenue to a framework for generalization by meta-learning, a technique that aims to generalize by learning how to benefit future learning from sampled episodes of related data or tasks (Hospedales, Antoniou, Micaelli, & Storkey, 2020). Previously, several methods achieved generalization by directly augmenting the training data using domain- and task-adversarial gradients (Zhou et al., 2022). How the proposed novel framework can algorithmically exploit the differentiability of the framework and control and the synthetic data generation of the preceding simulator for better generalization of the DDT is studied in this thesis.

Second, we focus on the integration of multi-system information, such as for the hydraulic and electrical subsystems of a hydro-power plant, using graph convolutional networks to enhance the generalizability and task performance of DDTs through a unified spectral-temporal representation. The primary objective is to develop a DDT that can learn complex dependencies in between different subsystems. While the planned framework for (multi-system) DDT generalization will be applicable to many different fields, we will develop the methodology for industrial systems and more specifically for power grids where we use combined electrical, hydraulic, and control system data. Industrial systems are optimal candidates to validate the research on DDT because we have access to conventional simulation, expert knowledge, and sensor data. Therefore, for our thesis work on digital twin generalization with meta and geometric deep learning, we select the case study of sensor fault and anomaly detection for railway traction networks.

2. RESEARCH QUESTIONS, WORK PACKAGES AND RESEARCH METHODS

The first objective of this thesis is to develop a framework that advances the generalizability of DDT. The secondary objective is to evaluate the proposed framework on the task of anomaly detection and sensor fault localization. The initial evaluation of this framework will be performed on simple differentiable simulators that are being developed in our current work (WP1). In the next step, we will extend its applicability to differentiable simulators of complex systems that can be trained both on real-world and on synthetic data (WP2). For this purpose, we will develop a graph neural simulator for power grid applications that is used to train the DDTs (WP3). For additional validation, we will use an existing neural simulator for battery degradation: Dynaformer (Biggio et al., 2022). Further to our methodological research on the differentiable simulator, we will design a novel DDT for railway traction networks in collaboration with Swiss Federal Railways (SBB). The goal is to improve the performance of state estimation and forecasting tasks in the Swiss railway traction network and to use the DDT to detect anomalies and sensor faults (WP4). In the following, we present the research questions (RQ) that we plan to address in this research as work packages (WP) from a top-down perspective.

2.1. WP1: A Meta-learning Framework for better Task Generalization

In WP1, we will develop a novel meta-learning framework that improves the generalization of an arbitrary DDT with the concept that we propose in the following. First, we will pretrain a model that we call learner with a user-defined task loss on synthetic data from a differentiable data source. Second, we will identify hard-to-process data (challenges) for the learner, based on task loss, and guide the training process of the learner to these challenges. While previous research has only optimized DDT generalization with respect to the user-defined task, we aim to optimize synthetic training data generation of a data-driven simulator with meta-learning. We will make use of the fact that we can generate additional challenging data with our synthetic data source and iterate in between the first two steps. Finally, we will be able to fine-tune the learner on real-world data to maximize task-specific performance. We, therefore, have to provide algorithms on two levels: on a meta-level, we will have to find points in the parameter space that are associated with challenging samples, and on the task level, we will have to optimize the learner. Therefore, we categorize this approach as meta-learning. We call the algorithm on the meta-level *ascent learning*, while we call the machine learning techniques for optimizing the learner weights *training*.

One of the main goals is to apply this meta-learning framework to industrial simulators. More specifically, the goal is to increase the generalization of DDT w.r.t slow-changing condition parameters such as parameters related to the degradation of components. However, the same strategy can also be applied more generally to increase the generalization of policy networks in reinforcement learning. The architecture of our proposed framework will combine several modules to achieve the generalization by meta-learning that we are aiming for.

In WP1, in the first step, we will perform a proof of concept for the proposed framework. This will include experiments on simple physics simulations such as a dampened pendulum or a pole cart system that is used to evaluate the performance in a reinforcement learning setting. Several challenges need to be addressed: firstly, we have to solve gradient compression to the manifold of physically viable configurations, if this is not addressed, new samples are not associated with valid operating conditions. Secondly, we have to increase data efficiency of the meta-learning process by optimized sampling of the operation conditions for the data generated during training. Thirdly, we have to tune the *ascent learning* hyper-parameters.

2.2. WP2: Meta-learning for DDT Generalization

In WP2, we will continue the conceptual work of WP1 and evaluate the performance of the proposed meta-learning framework on real-world applications. Moreover, we will extend the scope of our meta-learning framework to differentiable neural simulators and DDTs in the context of predictive maintenance. We aim to evaluate and improve the non-trivial *ascent learning* strategy in this application-oriented setting. That includes testing the impact of specific *ascent learning* strategies 1) on real-world task performance, 2) data-efficiency optimization for the DDT training, 3) verification of the stability of the user-defined task loss, and 4) the respective meta-learning hyper-parameter tuning Most importantly, we aim to demonstrate that our proposed *ascent learning* approach can improve the generalizability of a DDT. Furthermore, we aim to develop the necessary tools that are needed to analyze these increasingly complex applications.

In this work package, we intend to enhance the generalizability and performance of a variety of DDTs with our framework for real-world tasks: in one-step optimization, reinforcement learning, and for network resilience. We plan to use Dynaformer and our own neural simulator that is currently under development for the Swiss railway traction network. Based on the ongoing research in the XXXX lab and based on related literature, we will evaluate if there are any other neural simulators that would be suitable for the proposed framework. Furthermore, we plan to improve context generation for new operating conditions and context refinement with implicitly learned constraints in Noether networks (Antonova, Yang, Jatavallabhula, & Bohg, 2022).

2.3. WP3: Multi-system Information Fusion in Graph Convolutional Networks

For WP3, we propose to develop a neural simulator for combined state estimation and forecasting in schematic-based environments. We aim to develop a GCN algorithm that is able to achieve multi-system information fusion of heterogeneous data sources that consist of electrical, hydraulic, and control system data recorded from connected systems. The main novelty is information fusion in this environment. The main contribution of this WP will be a data-driven model that will learn the complex dependencies in this multi-system environment from data and will adopt task-specific data-driven DDTs from the power-grid state estimation and forecasting literature.

The proposed methodology will be based on a unified spectral-temporal latent space that represents the the dynamic state of the power grid. We aim to project fused multi-system information to this space to increase generalization and performance on all tasks. To find the optimal model for this representation, we will start with the architecture of a well-known GCN: StemGNN. We will explore different avenues to increase representative power and task performance: we will generalize the GCN's temporal graph representation, precondition and reinforce the attention mechanism on known, schema-derived graphs , and include physics-informed constraint (meta-)learning. Based on our current research, we expect that multi-system information fusion will not only improve the performance, but also improve the interpretation ability and robustness of a DDT.

In conjunction with WP4, we will validate the performance of our method on anomaly and sensor fault detection. The methodology will be evaluated on synthetic data generated in PowerFactory, Powerpanda, and from IEEE test feeders as well as on real data collected from the YYY traction network. If time permits, we will also evaluate the performance of the proposed framework on additional optional datasets to access the performance. These datasets are provided by ZZZZ and the Country's Meteorological Service.

2.4. WP4: Anomaly and Sensor-fault Detection & Isolation in Power Grids

In WP4, we plan to improve unsupervised anomaly detection and sensor fault isolation in power grids in terms of precision, recall, and evaluation time. The first scientific contribution is the exploration of the advantages of our proposed multi-system information fusion approach for anomaly detection. The second contribution is to narrow down the possible explanations for anomalies in a power grid. In our current research, we observe that the learned attention matrix in GCN for forecasting is sparse and aligns with the true hidden schema of the power grid. Hence, only a few local and relevant nodes are used for the forecast. We want to capitalize on this finding to improve sensor fault isolation. For the proposed methodology, we will use the simulator of WP3 and extend it with anomaly detection capabilities in a single DDT.

In practical applications, signals are usually acquired with a much higher sampling frequency than the technically possible update frequency of conventional state estimation. We aim to significantly reduce the gap in computation time to online real-time state estimation and anomaly detection because the DDT in planning is a data-driven neural network that lowers the computation barrier at inference time. This is also of particular interest to our industrial partner (YYY), because while the incoming signals are recorded every second in their system, the state estimation reference can only be updated every twenty seconds due to computational complexity. Currently, this leads to a considerable number of false alarms because the railway traction network is very dynamic and exceeds or falls below the reference value, even when the operation is safe. Additionally, this DDT will also be used for the framework developed in WP1 and WP3 as a practical case study. We will use a publicly available benchmark dataset to evaluate the results and will use, as a second case study, the real data from the YYY railway traction grid.

3. CURRENT PROGRESS

Research on WP1 has identified several key challenges that have been and are currently being addressed within the framework source code. We are currently evaluating the results across a wide range of datasets, including neural ODE-based differentiable simulators for problems in the physical domain, as well as GANs that address the image domain and are inherently differentiable. The early stages of the research on WP1 & 2 have been published as an extended abstract and presented orally at a conference. WP3 has concluded with a conference paper at PHME2024 that will be presented in July 2024. Several extensions to WP3 are either in discussion or currently being developed. One extension is set to be implemented as a master's thesis in Fall 2024. Based on the findings from WP3, current efforts are focused on integrating additional covariate information to improve the robustness and generalization of power grid load forecasting, which will also be published in June 2024.

ACKNOWLEDGMENTS

This research was funded by the Swiss Federal Office of Transport (FOT) under the project INtellIgenT maIntenance rAilway power sysTEms (INITIATE). The authors would like to thank FOT for the project coordination and Swiss Federal Railways (SBB) for providing the data for this research.

REFERENCES

Aivaliotis, P., Georgoulias, K., Arkouli, Z., & Makris, S. (2019, January). Methodology for enabling Digital Twin using advanced physics-based modelling in predictive maintenance. *Procedia CIRP*, *81*, 417–422. doi: 10.1016/j.procir.2019.03.072

Antonova, R., Yang, J., Jatavallabhula, K. M., & Bohg, J. (2022, June). *Rethinking Optimization with Differentiable Simulation from a Global Perspective* (No. arXiv:2207.00167). arXiv.

Biggio, L., Bendinelli, T., Kulkarni, C., & Fink, O. (2022, June). *Dynaformer: A Deep Learning Model for Ageing-aware Battery Discharge Prediction* (No. arXiv:2206.02555). arXiv.

Booyse, W., Wilke, D. N., & Heyns, S. (2020, June). Deep digital twins for detection, diagnostics and prognostics. *Mechanical Systems and Signal Processing*, *140*, 106612. doi: 10.1016/j.ymssp.2019.106612

Cuadra, L., Salcedo-Sanz, S., Del Ser, J., Jiménez-Fernández, S., & Geem, Z. W. (2015, September). A Critical Review of Robustness in Power Grids Using Complex Networks Concepts. *Energies*, *8*(9), 9211–9265. doi: 10.3390/en8099211

Hospedales, T., Antoniou, A., Micaelli, P., & Storkey, A. (2020, November). *Meta-Learning in Neural Networks: A Survey* (No. arXiv:2004.05439). arXiv. doi: 10.48550/arXiv.2004.05439

Hu, Y., Miao, X., Si, Y., Pan, E., & Zio, E. (2022, January). Prognostics and health management: A review from the perspectives of design, development and decision. *Reliability Engineering & System Safety*, *217*, 108063. doi: 10.1016/j.ress.2021.108063

Jardine, A. K. S., Lin, D., & Banjevic, D. (2006, October). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, *20*(7), 1483–1510. doi: 10.1016/j.ymssp.2005.09.012

Liao, W., Bak-Jensen, B., Radhakrishna Pillai, J., Wang, Y., & Wang, Y. (2022). A Review of Graph Neural Networks and Their Applications in Power Systems. *Journal of Modern Power Systems and Clean Energy*, *10*(2), 345–360. doi: 10.35833/MPCE.2021.000058

Soleimani, M., Campean, F., & Neagu, D. (2021). Diagnostics and prognostics for complex systems: A review of methods and challenges. *Quality and Reliability Engineering International*, *37*(8), 3746–3778. doi: 10.1002/qre.2947

Wang, J., Lan, C., Liu, C., Ouyang, Y., Qin, T., Lu, W., … Yu, P. (2022). Generalizing to Unseen Domains: A Survey on Domain Generalization. *IEEE Transactions on Knowledge and Data Engineering*, 1–1. doi: 10.1109/TKDE.2022.3178128

Zhou, K., Liu, Z., Qiao, Y., Xiang, T., & Loy, C. C. (2022). Domain Generalization: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–20. doi: 10.1109/TPAMI.2022.3195549