Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio

doi: 10.3850/978-981-18-8071-1_P473-cd

A Study on Gradient-based Meta-learning for Robust Deep Digital Twins

Raffael Theiler, Michele Viscione, Olga Fink
EPFL, Switzerland. E-mail: {first}.{last}@epfl.ch

Deep-learning-based digital twins (DDT) are a promising tool for data-driven system health management because they can be trained directly on operational data. A major challenge for efficient training however is that industrial datasets remain unlabeled. This is remedied by simulators that can generate specific run-to-failure trajectories of assets as training data, but extensive simulations are limited by their computational cost. Therefore, it remains difficult to train DDTs that generalize over a wide range of operational conditions. In this research, we propose a novel meta-learning framework that is able to efficiently generalize an arbitrary DDT using the output of a differentiable simulator. While previous generalization approaches are based on randomly-sampled data augmentations, we exploit the differentiability of the full pipeline to actively optimize the training data sampling by means of condition parameter's gradients. We use these gradients as an accurate tool to control the sampling distribution of the simulator, improving the representativeness, robustness, and training speed of the DDT. Moreover, this meta-learning approach leads to a higher quality of generalization and makes the DDT more robust to perturbations in the conditional parameters.

Keywords: Meta-learning, Deep Digital Twin, Differentiable Simulator, Digital Twin Generalization

1. Introduction & Motivation

DDTs are deep neural networks that encode the behavior of complex physical assets and have the ability to continuously adapt to operational changes. A major challenge in training DDTs is that their performance depends on the training data quality and representativeness Booyse et al. (2020). Manual estimation of the training distribution is non-trivial for prognostics and health management (PHM) because DDTs are required to work reliably under all operating conditions and must be robust to data distribution shifts. A recent trend in other applications has shown that DDTs can be successfully trained in purely simulated environments, which enables algorithmic tuning of the training sample distribution. There are two potential issues that should be noted. First, existing algorithmic tuning makes no use of the differentiability of the simulator Wang et al. (2022), and second, although an increasing number of differentiable simulators appear in the literature, they have little focus on improving DDT robustness.

2. Meta-learning Framework

In this research, we propose an approach that enables to increase the generalization and robustness of the DDT w.r.t operating condition parameters $\phi \in \Phi$ of a simulator. To achieve this, we first train the weights θ of a DDT with a user-defined task loss $\mathcal{L}$ on synthetic data $x \in \mathcal{X}$ from a differentiable data source, the differentiable simulator S. Second, we guide the training process of the learner to *challenging* (to learn) data samples $x^c \in \mathcal{X}$, identified by the loss and associated with ϕ^c, by gradient ascent and train the DDT on it. Finally, we alternate between the two steps until completion. We categorize this approach as meta-learning because on a meta-level, we shift the training distribution in Φ-space to increase the generalization efficiency of the training process while on the task level, we optimize the DDT to find optimal θ^*.

2.0.1. *Algorithm*

Given a differentiable simulator S, with initially sampled operation conditions Φ^T, and initially generated training data $\mathcal{X}^T = \{x_k \mid x_k = S(\phi_k^T)\}_{k=0}^P$, we find a sequence of loss-increasing condition parameters $[\phi_0^c, \ldots, \phi_i^c]$ starting from $\phi_0^c = \phi_k$ for i gradient ascent steps performed as:

$$\phi_{n+1}^c \leftarrow \phi_n^c + \beta \frac{\partial L_n}{\partial \phi_n^c} \tag{1}$$

We generate new *challenging* examples using the simulator $x^c = S(\phi_i^c)$. Finally, we optimize the DDT parameters θ using Adam optimizer.

3. Results

We consider the nonlinear damped pendulum as a standard benchmark case study for dynamical systems. The dynamical behavior of this model is governed by the ODE:

$$\ddot{\psi} + b\dot{\psi} + \frac{g}{L}\sin\psi = 0 \tag{2}$$

We implemented a video-generating S trained on the Runge-Kutta (RK4) integration of eq. 2. The DDT task is to estimate the frames of the video with pendulum velocity below a certain threshold (dead points). The $\mathcal{L}$ is the MSE between the estimated and actual dead points. In this scenario, ϕ represents the pendulum length L and b the damping factor. We compare our method to conventional training without ascent learning (baseline). Our approach with four gradient ascent steps (eq. 1) is demonstrated to be superior to the conventional method, as presented in Figures 3, 2, and 3. Since the data source is synthetic, both approaches have access to unlimited samples, hence both generalize. However, Fig. 3 shows that in our framework, the associated minimum mean loss is achieved more efficiently with 68% of the samples of the conventional method. Fig. 2 demonstrates a higher level of flatness of the loss surface, which corresponds to an overall increase in robustness and generalization of the DDT across various operating conditions Φ.

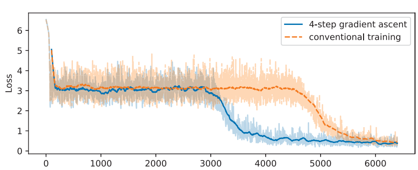

Fig. 1. Loss of conventional training compared to our method. We reach baseline generalization 32% faster.

4. Conclusion

Due to the high computational cost of simulators, the number of computations is often limited. Achieving a better generalization with the same computational budget is therefore desirable. In

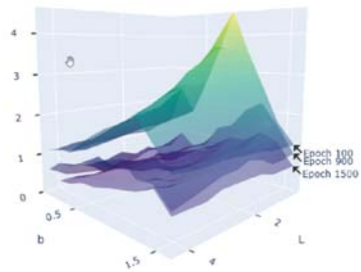

Fig. 2. Loss surface over the Φ space across epochs. The flatness of the loss landscape, including boundary regions, shows generalization and robustness across Φ-space.

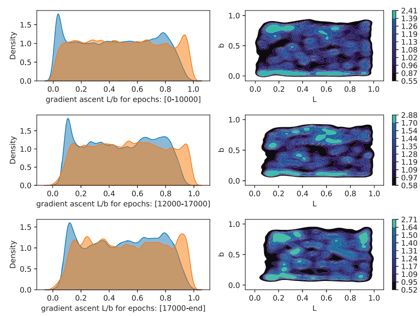

Fig. 3. Adaptive sampling densities in Φ at the beginning, middle, and end of training. L in orange and b in blue (of eq 2).

such cases, our framework is an attractive choice for data-driven safety and reliability applications as it improves robustness and sampling efficiency while increasing DDT task performance and retaining fully trainable components. The achieved flatness of the loss landscape implies robustness with respect to condition parameter perturbations, making the DDT more resilient. By exploiting the increasing availability of end-to-end differentiable industrial simulators, our meta-learning framework is applicable to a growing set of applications for DDTs.

References

Booyse, W. et al. (2020, June). Deep digital twins for detection, diagnostics and prognostics. *Mechanical Systems and Signal Processing 140*, 106612.

Wang et al. (2022). Generalizing to Unseen Domains: A Survey on Domain Generalization. *IEEE Transactions on Knowledge and Data Engineering*.