Towards a Digital Twin for Underwater Systems Based on Meta-Learning

Marina Rantanen Modéer
Saab Dynamics AB
Bröderna Ugglas gata
SE-581 88 Linköping
SWEDEN

marina.rantanenmodeer@saabgroup.com

ABSTRACT

This paper presents the initial research to achieve a digital twin based on the concept of meta-learning. The primary objective of these first steps towards building the desired digital twin is to evaluate appropriate methods befitting the circumstances of defence industry. The main challenge is how to train and validate a model when data is scarce - a common theme in the defence sector. The use case at hand is a lithium-ion battery employed in an advanced underwater defence system, aiming to create a digital artefact that may be employed for both maintenance and design purposes. The initial implementation uses Recurrent Neural Networks (RNN) and Model-Agnostic Meta-Learning (MAML) to implement an inner and outer learning loop to achieve a learner that can adapt to new tasks quickly. The base learner leverages an open data set of battery degradation to generate gradients for the meta-model. The results underscore the digital twin's potential as a valuable tool for informed decision-making, fostering reliability, and readiness in underwater defence systems.

1.0 INTRODUCTION

In late 2022, Saab delivered the state-of-the-art torpedo system Saab Light Weight Torpedo (SLWT), illustrated in Figure 1, to the Swedish Defence Materiel Administration (FMV) and the Swedish Armed Forces. The SLWT is roughly 2.85 m long with a diameter of 0.4 m and a total mass of about 340 kg. The propulsion system of the SLWT integrates a pump jet drive, an electric DC motor with a gearbox and a rechargeable lithium-ion battery.

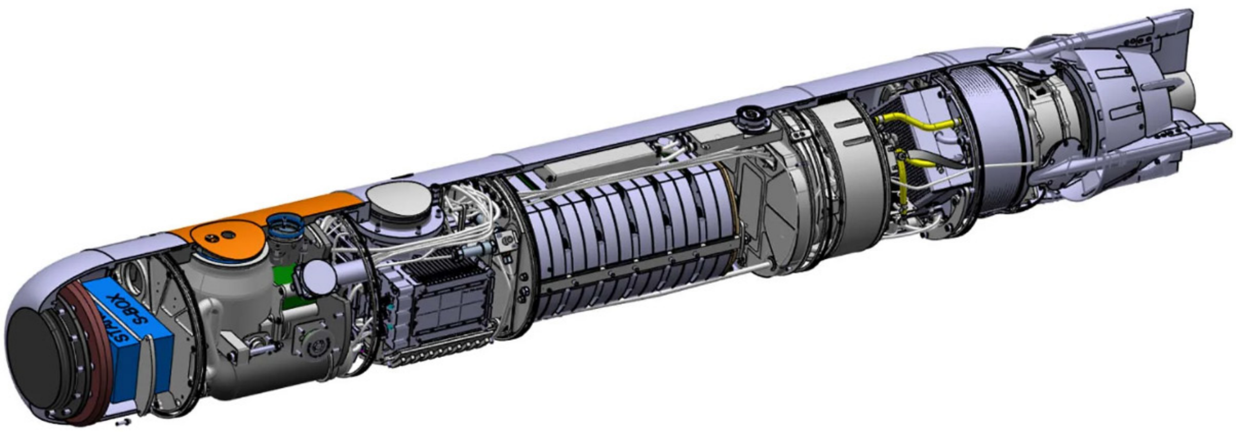

Figure 1: Design illustration of the Saab Light-Weight Torpedo (SLWT), sourced from FMV [15].

As such, the design sets out to maximize operational safety and system availability and minimize the number of components at risk of corrosion and other types of detritions as opposed to other design-options based on e.g., combustion engines. It is worth noting that autonomous and other automated underwater systems are challenging to design - not only because of complex behaviour in a communication contested environment, but also due to the unhospitable traits of the sea, which suggest a particular class of physical and technological complexity. Proper testing is crucial to mitigate potentially expensive design flaws, but full-scale models are not realistic to build until the design is more or less finalized due to resource limitations. The choice of the propulsion system of the SLWT was naturally made with respect to stakeholder requirements. However, batteries inevitably degrade over time and the design as well as the maintenance plan of the system should ideally integrate detailed knowledge about the rate of degradation to optimize system performance throughout its lifespan. A Digital Twin (DT) is therefore being developed through a simulation model that is built using Machine Learning (ML) techniques for each unit of a selected number of SLWT entities, which serves two ultimate purposes:

1) Support data-driven, optimal operation of the SLWT through predictive maintenance of (the propulsion system of) existing torpedoes

2) Support data-driven design of next generation underwater defence systems

The goal is threefold in the first stage of the development of the DT: enable effective prediction of battery degradation for maintenance, produce accurate synthetic data through simulations, i.e., alternatives to empirical data, for determining the optimal operational envelope of the SLWT energy module and create reliable models of various energy aspects for future design. Model-based optimization of the operational phase using data-driven representations of the already existing assets of the SLWT (i.e., a DT) could help diagnose the assets for possible faults and may drastically cut costs and risks of errors in the future.

To tackle the construction of the DT, the work presented here deals with to separate themes - first, the issue of designing and validating a satisfying representation of the battery of the SLWT and second, the challenge of data access, since the amount and quality of data largely defines which ML techniques can achieve a satisfying model of the battery in question.

2.0 DIGITAL TWINS IN INDUSTRY

Multiple sources suggest that DTs have emerged as a potentially powerful tool in industrial settings and "opens up a world of possibilities" [5], enabling companies representing practically the entire engineering spectrum to create virtual representations of physical assets and systems. These virtual counterparts are built to provide real-time insights, predictive analytics, and simulation capabilities beyond traditional methods. There are several definitions of DTs, one of which is provided by the Defense Acquisition University and is formulated as follows:

> *A virtual replica of a physical entity that is synchronized across time. Digital twins exist to replicate configuration, performance, or history of a system. Two primary sub-categories of digital twin are digital instance and digital prototype.* [6]

This section explores the development and use of DTs in defence applications from an industrial perspective, highlighting their benefits and impact.

2.1 Benefit of Digital Twins in Industrial Settings

The application of DTs in industrial settings has revolutionized various industries, offering numerous benefits and opportunities for optimization and improvement, some of which are illustrated in Figure 2.

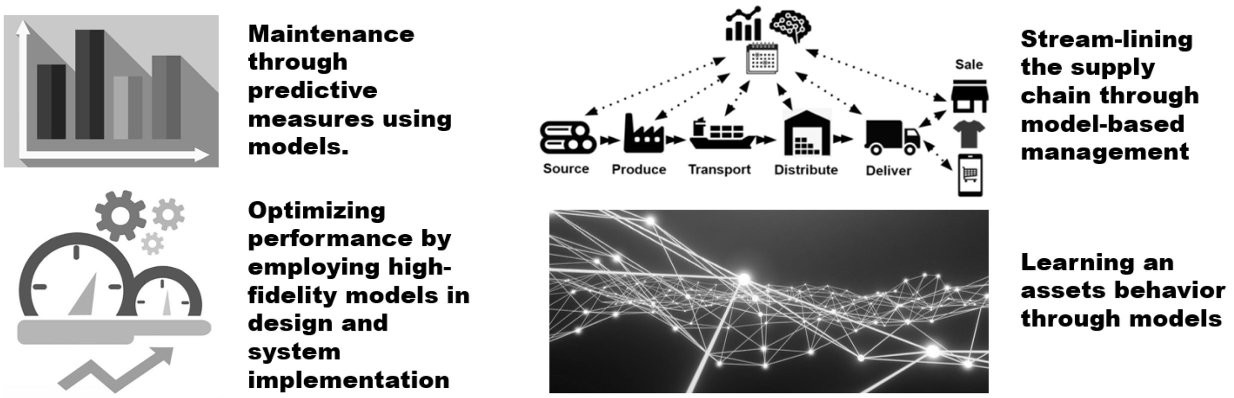

Figure 2: Benefits of employing digital twins in industry, e.g., predictive maintenance, performance optimization, training and simulation, and supply chain management.

DTs enable predictive maintenance strategies by continuously monitoring – both based on recorded data and based on extrapolation through generated data - the asset's performance and detecting early signs of potential failures. Secondly, they provide insights into the asset's performance and efficiency, as engineers can simulate different operating conditions, test optimization strategies, and identify areas for improvement. Furthermore, DTs serve as valuable tools for training and simulation purposes. Operators and maintenance personnel can interact with the virtual model, gaining hands-on experience and familiarity with the asset's behaviour. DTs additionally offer visibility into the entire supply chain, enabling companies to track and optimize inventory levels, monitor logistics operations, and simulate scenarios for demand forecasting. Finally, DTs support the entire lifecycle of an asset, from design and manufacturing to operation and decommissioning. By capturing data and insights throughout the asset's lifecycle, companies can make informed decisions, optimize asset utilization, and plan for replacements or upgrades effectively, as suggested by Figure 3.

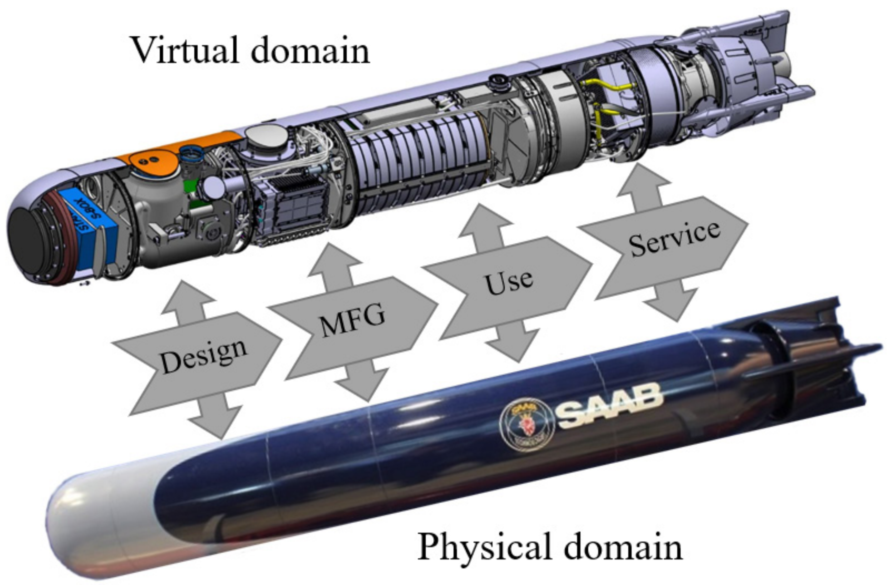

Figure 3: A digital twin ideally supports the entire lifecycle of a system throughout design, manufacturing, use and service/maintenance.

2.1.1 Benefits and Challenges in Defence Industries

The challenges of data infrastructure for DTs in the defence industry are multifaceted and demand careful consideration. Data collection and integration present significant hurdles. DTs rely on vast amounts of heterogeneous data from various sensors, platforms, and systems. Integrating and harmonizing this data in a standardized format is a complex task, especially when dealing with legacy systems and classified information, which may have different protocols and security requirements. By definition, a DT must be built using data recorded from a physical asset, which links to the fundamental research question of this paper. Data is clearly vital in the age of ML, but aerospace and defence (amongst others) require highly developed procedures for data protection, why assurances of data privacy and security have become a critical obligation as part of ML development. There are obvious concerns of the risks to personal and proprietary data leakage, misuse, and abuse, especially for cloud-based and other distributed solutions where fundamental infrastructure is in the hands of a third party. For building DTs of defence systems, ML techniques must be chosen with respect to limited data access and strict data handling procedures. Ensuring data security and confidentiality is paramount in the defence sector. Protecting sensitive information from unauthorized access or cyber threats is crucial, considering the potential implications of data breaches in military operations. Implementing robust encryption, access controls, and secure data transfer mechanisms becomes essential to maintain the integrity and confidentiality of the DT's data infrastructure.

Another central question is that of scalability and real-time data processing, as they pose significant challenges. Defence applications often demand real-time decision-making, necessitating data infrastructures that can handle high volumes of data and process it rapidly to support time-sensitive operations.

It is also evident that data interoperability remains a challenge when collaborating with different defence agencies or allied nations, such as within NATO. Establishing common data standards and protocols is crucial to ensure seamless data exchange and effective interoperability between DT systems. Addressing these challenges requires a holistic approach, involving collaboration between defence organizations, technology providers, and policymakers to create robust and adaptive data infrastructures for digital twins in the defence industry.

In the context of this paper, the primary purpose of the desired DT is to enable a tool for real-time monitoring and predictive maintenance of the selected SLWT assets. By simulating the behaviour of their batteries based on a high-fidelity model of each individual, the idea is to gain insights into health, performance, and RUL aspects of each asset's battery. Optimization and performance enhancement are expected to be facilitated through exploring various operating conditions and configurations. This in turn is expected to allow for fine-tuning battery parameters to achieve optimal performance under different scenarios in the long run - especially for future designs and other system concepts. Additionally, the DT is projected to aid in risk assessment and mitigation. By subjecting the virtual battery to simulated extreme conditions and failure scenarios, vulnerabilities can potentially be identified and addressed, leading to improved reliability and safety of the entire SLWT. As such, the DT is expected to foster a better understanding of the battery's interactions with other components in the system, leading to improved system integration and overall efficiency.

2.2 Development of Digital Twins

The development of DTs involves the integration of various technologies, including sensors, Internet of Things (IoT) devices, data analytics, and simulation models, potentially expressed in a multitude of languages and tools. These technologies work together to capture, monitor, and analyse real-time data from physical assets and feed it into the DT model.

As a first step, simulation models are developed based on the selected physical asset's characteristics and operational parameters. These models replicate the behaviour and dynamics of the asset, allowing engineers to analyse and predict its performance under various conditions. Typically, sensors and IoT devices are deployed to collect data from the physical assets. These devices capture the information of interest (such as temperature, pressure, vibration, operational parameters and the like). The data is then transmitted to the DT for analysis and further model fine-tuning. The collected data is normally integrated into a unified platform where it is processed and combined with other relevant data sources, such as historical records, maintenance logs, and supply chain information. This integration enables a comprehensive view of the asset's behaviour and performance.

Once the DT is developed and deployed, it should ideally receive continuous real-time data from the physical asset. However, in defence related settings, data transfer needs special attention as defence secrecy and system security are key mechanisms. This subject will therefore be explored in greater detail below. Regardless of the data infrastructure, the data is compared with the simulation models, enabling engineers to monitor and identify any deviations or anomalies in the asset's behaviour.

Another vital area is the abstraction level of the model in question. A DT model normally would abstract the essential features and dynamics of the physical system while omitting unnecessary intricacies. Representing the system at too granular a level may lead to excessive computational overhead and make it challenging to use the digital twin for real-time applications. Conversely, too high-level abstraction might sacrifice accuracy and overlook critical behaviors. Therefore, the ideal abstraction level for a DT should capture the relevant physical processes, interactions, and dependencies while ensuring efficient simulation and responsiveness for real-time monitoring, control, and decision-making tasks. Regular validation and calibration against the physical system are essential to ensure that the digital twin remains an accurate and reliable representation.

2.2.1 Relevant Artificial Intelligence Methods

There are several Artificial Intelligence (AI) methods that may be employed for training a digital twin. ML algorithms, such as support vector machines (SVM), random forests (RF), and deep learning networks (DLNs), comprise some of the typical tools for data-driven modelling and behaviour prediction that appear in literature. Additionally, hybrid approaches that integrate multiple AI methods could potentially provide a comprehensive framework for digital twin development, as suggested in Ref. [9]. For the purpose of building battery degradation models, these AI methods show thought-provoking potential. Several sources indicate that they enable state-of-charge (SoC) estimation, Remaining Useful Life (RUL) prediction, fault diagnosis, optimal charging and discharging strategies, and adaptive control, which would contribute to enhancing the overall performance and management of battery systems, as suggested in Refs. [10], [11] and [12] among others. However, the prerequisite here is that data is scarce or simply confidential, which rules some of the suggested methods out. Instead, Federated Learning (FL) appears in recent literature as an alternative to methods demanding direct access to data, as suggested in Ref. [13]. It has emerged as a prospective solution that facilitates distributed collaborative learning without disclosing original training data and therefore shows some potential. Unfortunately, retaining data and computation on-device as in FL are not sufficient for privacy-guarantee because model parameters exchanged among participants conceal sensitive information that can be exploited in privacy attacks. So, going back to square one. It is well-known that models that are built using various ML techniques are usually more accurate the larger the amount of data that they are based on is, as suggested in Refs. [2] and [13] among many others. However, human learning processes – in contrast to most machine learning processes – are usually exceptionally effective despite being exposed to small quantities of data. How can a machine do what humans do and learn *how* to learn a task, rather than just learning a task?

Meta-learning is about training a "learner", i.e., a model that "learns how to learn". The concept consists of a pretraining stage where a model may be trained on ample amounts of labelled data that can be sourced from another domain, task, or provider, which is then followed by a fine-tuning stage where the model is further trained in the target domain on the task of interest. It may also enable so called "few-shot learning", i.e., learning given only a few examples (called shots), in the target domain, in which a classifier must adapt to distinguish novel classes not seen during training of these classes [4]. The concept is similar to transfer learning, though there are some distinctions to be made. Meta-learning and transfer learning are both strategies employed in ML to improve the performance and generalization of models. The fundamental similarity between these two approaches lies in their focus on leveraging knowledge or experience gained from one task or domain to enhance performance on another, as mentioned above. In meta-learning, the emphasis is on learning how to learn, where the model acquires higher-order knowledge across multiple learning tasks, enabling it to adapt in a relatively short amount of time and efficiently to new, unseen tasks. Transfer learning, on the other hand, involves using pre-trained models on a source task and fine-tuning them on a target task, allowing the model to transfer its knowledge from the source domain to the target domain, as already suggested and discussed in Ref. [4]. It is worth noting that the major difference is that transfer learning expects that tasks are mostly similar to each other, but meta learning does not. In transfer learning, model parameters are pre-trained with a large dataset. Those parameters are then used as initial parameters to finetune on some other tasks having a smaller dataset. This classic pre-training approach has no guarantee of learning an initialization that is good for fine-tuning. In meta-learning, an initial set of parameters are learned that can be finetuned relatively easily on another similar task with only a few gradient steps. It directly optimizes performance with respect to this initialization differentiating through the fine-tuning process. Both meta-learning and transfer learning thus harness the benefits of reusing learned representations, with the aim to create more robust and effective learning paradigms that can excel even in scenarios with limited data or novel tasks.

When it comes to training a meta-learner for a digital twin of a battery, few AI methods stand out as particularly applicable. A possible structure is that of a combination of an inner loop that trains a base model on specific tasks and an outer loop that takes gradients as inputs from the inner loop to observe the *how* of the training. For the inner loop two promising approaches are Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), as illustrated in Figure 4. They are suggested by current literature specifically for extracting spatial and temporal patterns from battery data [7], [8]. On top of that, literature suggest Model-Agnostic Meta-Learning (MAML) for the outer loop as a method that excels in scenarios with limited data for new tasks, making it a natural fit for the challenges posed by battery behaviour prediction, as e.g., suggested in Ref. [14].

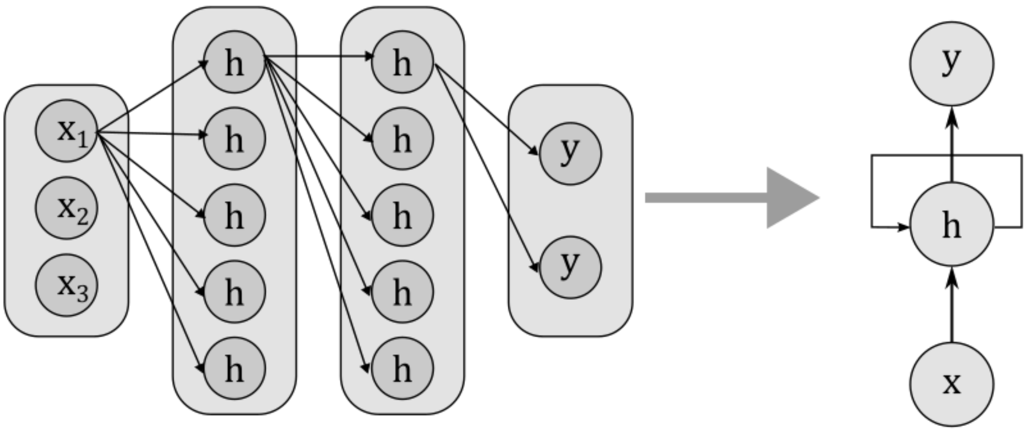

Figure 4: Recurrent Neural Network (RNN) using a set of inputs (x) and hidden layers (h) to generate an output (y).

Hence, MAML and RNN, with memory-augmented architectures or gradient-based meta-learning methods, are of interest for this work for enabling the meta-learner to adapt and generalize to unseen battery systems or varying operating conditions by leveraging knowledge from previously observed batteries, in contrast to other AI methods that require a large amount of data. The ultimate purpose of the meta-learner is to effectively learn to adapt integrated parameters and strategies based on the characteristics and behaviour of different batteries, enhancing the general ability to make accurate predictions. MAML and RNN are central concepts for the work presented here and will be discussed in further detail below, with a suggestion of implementation and discussion thereof.

3.0 TOWARDS A DIGITAL TWIN OF THE SLWT

The DT presented in this study is considered a work-in-progress and as such, the focus of this section is on showcasing the results achieved thus far. The process of building an effective DT using MAML with the base learner trained on the Oxford Battery Degradation Dataset is ongoing and subject to further refinement. Nevertheless, the attained outcomes offer valuable insights into the potential of the meta-learner in predicting battery behaviour. This section presents the fundamental ideas for the DT design and further sections will present current findings, demonstrating the feasibility and promising prospects of the DT approach for optimizing battery management and performance prediction in real-world applications.

3.1 Relevant First Principles' Models

Battery modeling is a challenging task. The central first principles models of a lithium-ion battery for modeling state-of-health (SoH) are based on fundamental electrochemical principles and physical processes that occur within the battery. These models are more detailed and physics-based compared to empirical or data-driven models and are naturally of interest for battery design. The following summarizes some of the central first principles models for modeling the state-of-health of a lithium-ion battery to capture what information the ML techniques of choice are meant to model. Three of the main modeling principles comprise:

1) ***Single Particle Model (SPM):*** The SPM is a widely used model for modeling the electrochemical processes within a single electrode particle of a lithium-ion battery. It considers diffusion of lithium ions, solid-phase diffusion, and electrochemical reactions within the particle. SPM can provide insight into aging mechanisms related to solid electrolyte interface (SEI) formation and growth.
2) ***Electrochemical Thermal Model (ETM):*** The ETM combines electrochemical and thermal modeling to predict battery performance and thermal behavior simultaneously. It takes into account heat generation during charge/discharge cycles, which affects the battery's aging and SoH.
3) ***Doyle-Fuller-Newman (DFN) Model:*** The DFN model is a multi-dimensional, spatially resolved model that considers the coupling between electrochemical processes and transport phenomena (e.g., diffusion, migration) in the battery electrodes. This model can help analyze the effects of electrode degradation and impedance growth on battery SoH.

These first principles models can provide valuable insights into the aging mechanisms and state-of-health of lithium-ion batteries under different operating conditions. However, they are computationally intensive and may require parameter estimation and calibration using experimental data for accurate predictions. In some cases, simplified versions, or reduced-order models of these first principles models are used for faster computations while capturing essential aging effects. Various electro-chemical models are being developed within the scope on this research as they are fundamental to the understanding of battery design. Some of the electrical and electro-thermal models of battery degradation that are the result of the meta-learner developed by this work will be based on hybrid methods of measured, empirical data and data sourced from fundamental models.

3.1.1 Model Abstraction Level

The suitable abstraction level of the model meant for the battery DT is a crucial consideration to achieve an accurate representation of the real-world SLWT while maintaining computational efficiency. Striking the right balance between complexity and simplicity is key. Battery research has emerged as a highly significant subject area and there are yet many crucial components missing in its body of knowledge when it comes to modeling, especially on how to predict battery performance over time. Oxford sources conclude:

> *Time, usage, and environmental conditions lead to performance deterioration and cell failures […] The physical and chemical mechanisms responsible for degradation are numerous, complex, and interdependent. Our understanding of degradation and failure of Li-ion cells is still very limited and more limited yet are reliable and practical methods for the detection and prediction of these.* [1]

As the goal here is just that - predicting battery degradation in a precise and reliable manner – the DT must represent the suggested physical and chemical mechanisms of the real asset on a suitable level of abstraction. In the context of this work, the aim is to calculate State of Health (SOH), i.e., battery performance relative to original capacity, as follows:

$$SOH = \frac{C_{FULLCHARGE}}{C_{DESIGN}} (\%)$$

There is a distinct degradation curve of the SOH with the charging cycle number, i.e., the number of times the battery has been charged and discharged. The Oxford Dataset [1] contains extensive SOH data of eight different li-ion cells for a given laboratory experiment and is therefore considered useful in the context of this work. Figure 5 shows the plots of temperature against voltage of the first charging cycle of each of the eight cells in the Oxford Dataset.

As a first step, this work has aimed at training a model that can predict the temperature of any new li-ion cell based on voltages as input data, in accordance with the plots in Figure 5.

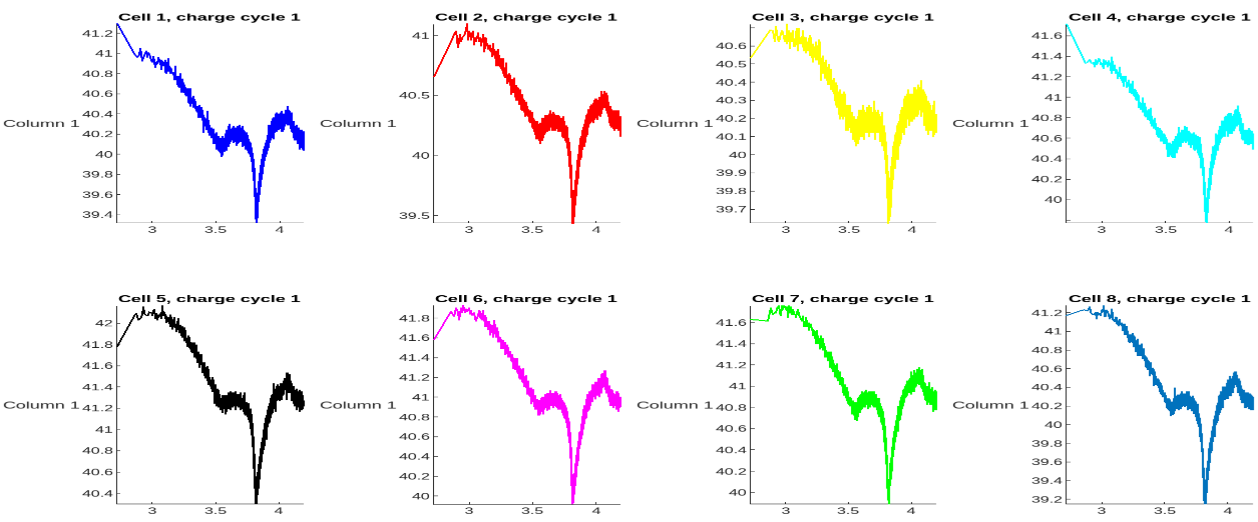

Figure 5: Temperatures plotted against voltages of the eight li-ion cells in the Oxford Dataset.

3.3 Meta-Learning Setup – the Base Learner and Outer Learner

At this point it is suitable to remind the reader that this DT is meant to be trained on small amounts of data. The suggest “base learner” and the “outer learner” have already been established as components with distinct roles above, here illustrated in Figure 6. The base learner is responsible for learning from the data within each task or individual learning episode. It represents the learner that adapts to specific examples and makes predictions within a given task context. The base learner's objective is to quickly learn and generalize from the available data in order to perform well on the specific task it is presented with. It typically undergoes multiple iterations of learning on different tasks or episodes during the meta-training phase.

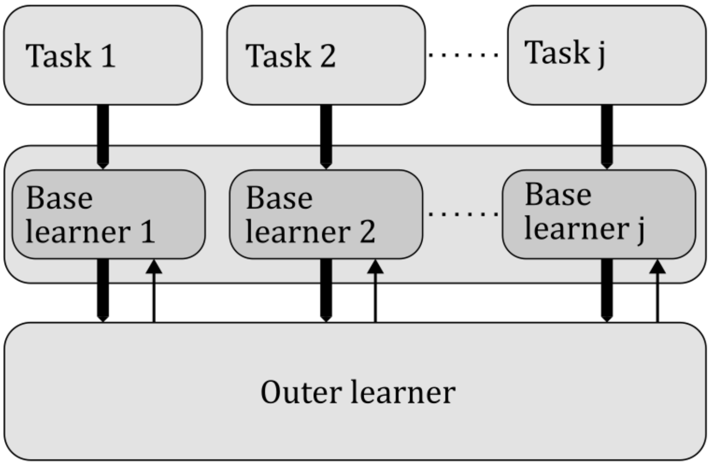

Figure 6: A meta-leaner built on a set of base learners. The use case presented in this paper uses RNNs for the base learner and MAML for the outer learner.

On the other hand, the outer learner, operates at a higher level and takes into account the knowledge and experience gained across multiple tasks or learning episodes. Its purpose is to learn from the base learners' performance and extract general patterns or meta-knowledge that can facilitate efficient learning and adaptation across new tasks. The meta-learner's objective is to discover effective initialization, optimization, or update strategies for the base learner. It can be viewed as a "learning to learn" algorithm that guides the learning process of the base learner.

In more detail, the process starts by the base learner solving each task in a set of j tasks. Then the base learner is typically split into two parts that correspond to 1) task-specific parameters and 2) meta-parameters. The base learner subsequently updates task-specific parameters during adaptation to different tasks. This is followed by the outer learner (i.e., the meta-learner) accumulating experiences from multiple tasks, mines their shared features, minimizes the associated loss function on validation data to minimize the generalization error and updates the meta-parameters in the base learner. The base learner also learns task-specific features by updating task-specific parameters during adaptation to each unseen task. The meta-learner models features shared by all tasks, guides the base learner to better generalize to unseen tasks, and maximizes generalization capability of the base learner. The base learner supplies task-specific information and a generalization error on task validation data to the meta-learner. The meta-learner aggregates task-specific information to learn features shared by all tasks and minimizes generalization error of the base learner to update the meta-parameters. The meta-learner supplies updated meta-parameters to the base learner. Communication between the base learner and the meta-learner may be setup such as to facilitate closer cooperation between the base layer and the meta layer [16].

4.0 META-LEARNER BASED ON OPEN DATA SETS

By harnessing the potential of MAML, the aim here is to develop a meta-learner capable of quickly adapting to new tasks, to make it a candidate for battery management. The following section outlines the essential steps involved in building the first version of the digital twin, with the base learner based on the Oxford Battery Degradation Dataset, sourced from Ref. [1].

4.1 Training a Meta-Learner for an SLWT DT

4.1.1 Training Architecture

In accordance with the above-described rationale, the first meta-learner has been designed with a base learner constructed as an RNN, with a Long Short-Term Memory (LSTM) structure and an outer learner with a Model-Agnostic Meta-Learning (MAML) structure to update the base learner's parameters such that it can quickly adapt to new tasks, as illustrated in Figure 7. The base learner takes voltage and temperature data from battery cell 1 as provided by the Oxford Battery Degradation Dataset and referenced in Ref. [1] to try to predict the behaviour of the corresponding parameters of another battery cell in the same set. The design is set up such that meta-learner learns to update the base learner model's parameters in a way that it has the potential to adjust to new tasks with limited data. Additionally, a meta-testing algorithm has been designed to assess the trained meta-learner's ability to generalize and make accurate predictions for new tasks without extensive retraining.

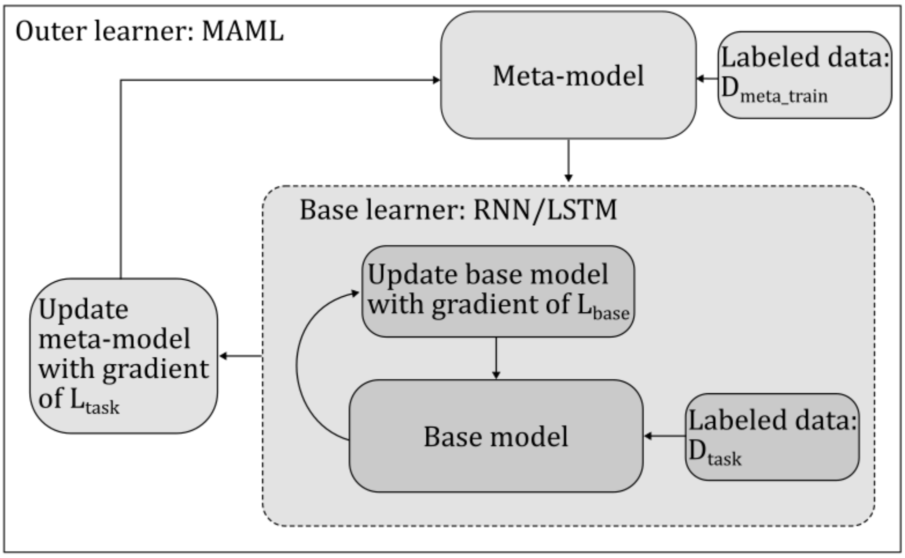

Figure 7: Meta-model architecture: Model-agnostic meta-learning and recurrent neural network.

An algorithm corresponding to the MAML and RNN architecture illustrated in Figure 7 has been designed and is outlined in Algorithm 1 and 2 below. The implementation of the specific loss functions, gradient computations, and update strategies for the chosen base learner and meta-learning algorithm are of essence. Also, hyperparameter tuning and cross-validation may be used to further optimize the performance of the meta-learner model.

Algorithm 1: Model-agnostic meta-learning

Require: $p(T)$: distribution over tasks

Require: α, ß: step size hyper parameters

1: randomly initialize θ

2: while not done do

3: Sample batch of tasks $Ti \sim p(T)$

4: for all Ti do

5: Evaluate ∇θ L Ti (*f*θ) with respect to K examples

6: Compute adapted parameters with gradient descent: θ′i = θ – α ∇θ L Ti (*f*θ)

7: End for

8: Update θ ← θ − ß ∇θ Σ Ti ~ p(T) L Ti (*f*θ′)

9: end while

Algorithm 2: Adaptation MAML

1: function ADAPT (*f*, ϴ, Da; ϕ)

2: ϴ0 ← ϴ

3: for j ϵ { 1 … adaptation steps } do

4: Lj ← L (Ya, *f* (Xa; ϴ j-1))

5: ϴj ← ϴ j-1 - ϕ ∇ ϴ j-1 L j

6: return ϴ adaptation steps

The suggested algorithm, as presented by Algorithm 1 and 2, have been formulated in MATLAB using the Deep Learning Toolbox, provided by Mathworks.

4.1.2 Results from the Trainer

The final adapted base learner's predictions are plotted against the original data to visualize the model's effectiveness in predicting battery temperature, as presented in Figure 8 and Figure 9. Figure 8 displays the training result of 750 epochs, where data has been split such that the first 80% were used for training and the last 10% were used for validation and test, respectively.

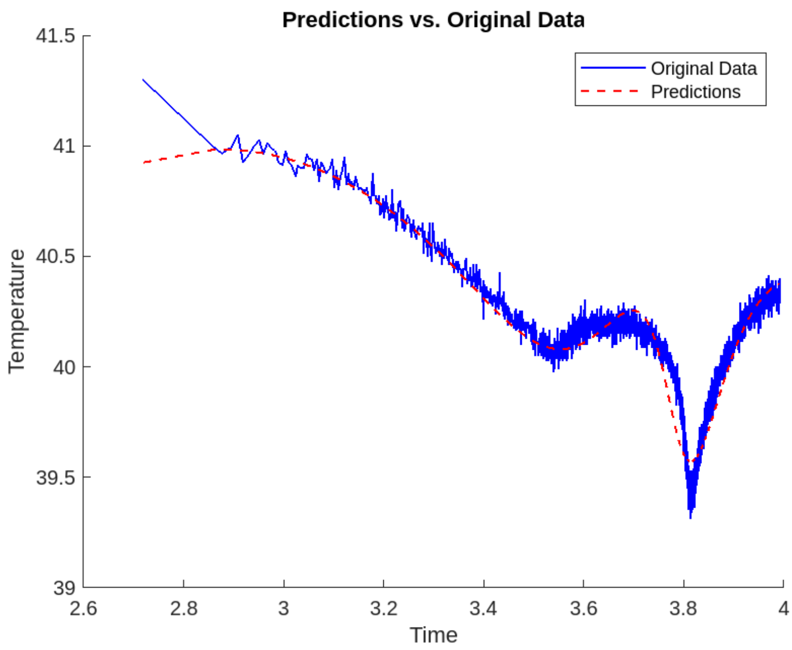

Figure 8: Training a model of temperature prediction based on voltages from one cell.

The training progress of the base learner is illustrated in Figure 9 and the predicted temperatures of the resulting model are plotted against the actual temperatures, as illustrated in Figure 10. Furthermore, a histogram of the errors of the prediction model is shown in Figure 11.

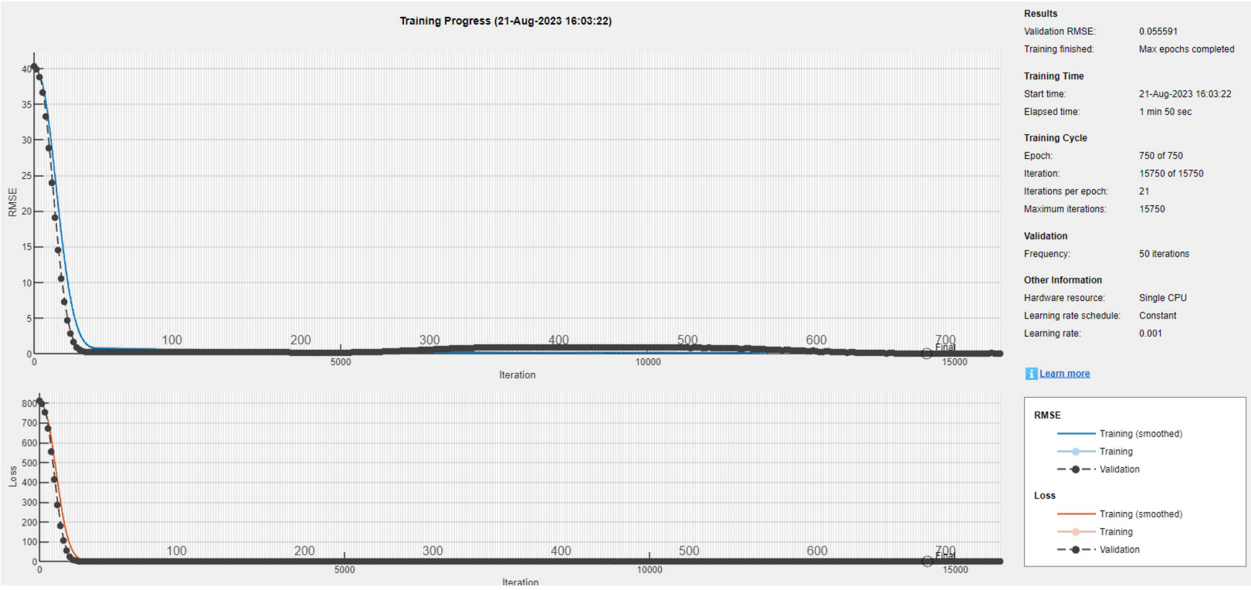

Figure 9: Training progress of the presented RNN.

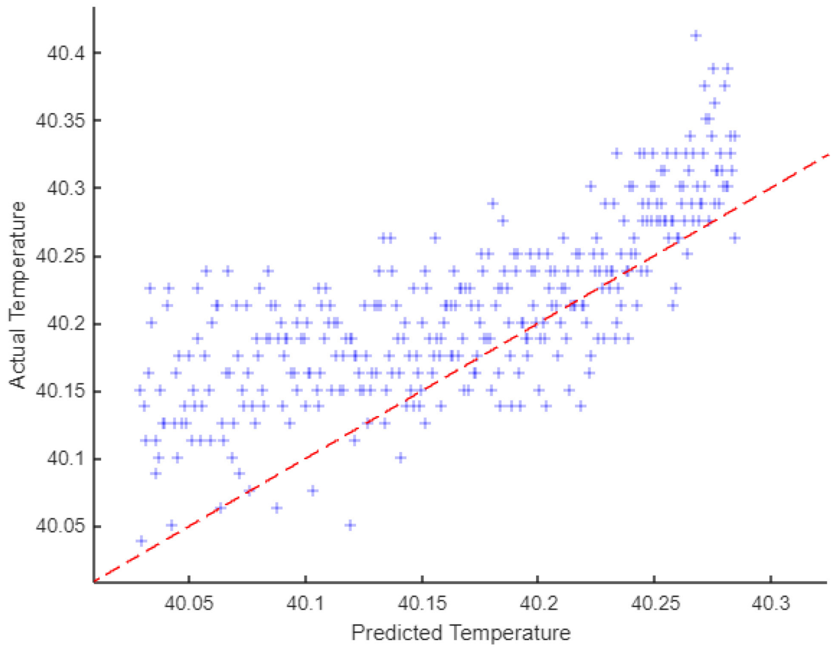

Figure 10: Predicted temperatures plotted against actual temperatures.

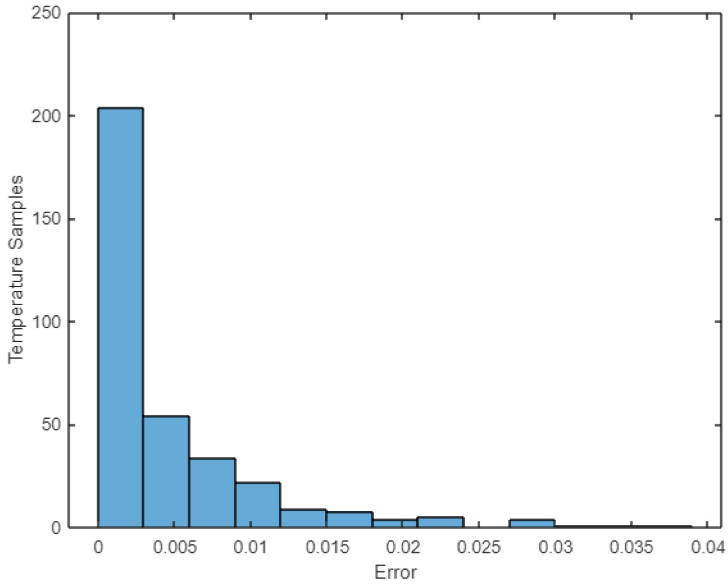

Figure 11: Errors of temperature prediction model.

The proposed algorithm successfully demonstrates the use of MAML with an RNN base learner implemented with an LSTM structure for the prediction of thermal development in li-ion batteries given voltages as input, an initial learning rate of 0.001 and 750 epochs. Experimental results suggest that the meta-learner adapts to new tasks with limited data, making it a promising approach for real-world battery management systems. However, the model's predictions may lose accuracy due to several factors. First, the complexity of the battery system and its intricate behaviour might not be fully captured by the chosen base learner. This could result in the model struggling to generalize well to unseen data and various operating conditions. Moreover, the training process of the MAML-based meta-learner might encounter challenges in balancing exploration and exploitation, resulting in suboptimal adaptations during the inner loop updates.

5.0 CONCLUSIONS

In the defence sector data may be challenging to access, both due to the fact that some systems are rarely used in the field and due to high data security standards. This makes data-driven modelling techniques challenging to employ and digital twins may be difficult to build as they typically require large quantities and continuous streams of data to improve their accuracy. Therefore, the ongoing research presented here investigates how to build high-fidelity models of real-world assets using small data.

The idea builds on training a base-learner on ample amounts of open data and then training a meta-learner to learn how to learn new tasks based on observations made during the training of the base-learner. The intermediate results in this work presents a setup based on Recurrent Neural Networks (RNN) using open datasets to generate gradients from an inner loop of a Model Agnostic Meta-Learner (MAML). The suggested algorithm is implemented in MATLAB using the Deep Learning toolbox, which shows some promise. The MAML setup creates a meta-model on which to train the desired target model in the target domain. This has been built to constitute an initial digital twin of a lithium-ion battery, whose predictive capabilities hold significant potential for optimizing battery performance and prolonging its lifespan in various real-world scenarios. This first version of the battery DT will be used for further development.

The ability to support data-driven decision-making allows for continuous improvement and adaptation based on real-world performance feedback, ensuring that the future development of underwater defence systems at

Saab remains at the forefront of technology and operational effectiveness. Future work may explore extending the approach to other types of degradation prediction and benchmarking against alternative meta-learning algorithms.

6.0 REFRENCES

[1] https://ora.ox.ac.uk/objects/uuid:03ba4b01-cfed-46d3-9b1a-7d4a7bdf6fac.

[2] Zhu, Xiangxin, et al. "Do We Need More Training Data or Better Models for Object Detection?." BMVC. Vol. 3. No. 5. 2012.

[3] Li, Li, et al. "A review of applications in federated learning." Computers & Industrial Engineering 149 (2020): 106854.

[4] Gevaert, Olivier. "Meta-learning reduces the amount of data needed to build AI models in oncology." British Journal of Cancer 125.3 (2021): 309-310.

[5] Javaid, Mohd, and Abid Haleem. "Digital Twin applications toward Industry 4.0: A Review." Cognitive Robotics (2023).

[6] https://www.dau.edu/glossary/Pages/GlossaryContent.aspx?itemid=27349

[7] Schmitt, Jakob, Ivo Horstkötter, and Bernard Bäker. "Electrical lithium-ion battery models based on recurrent neural networks: A holistic approach."Journal of Energy Storage 58 (2023): 106461.

[8] Wang, Shunli, et al. "A critical review of improved deep convolutional neural network for multi-timescale state prediction of lithium-ion batteries." Energies 15.14 (2022): 5053.

[9] Gao, Xinlei, et al. "Designed high-performance lithium-ion battery electrodes using a novel hybrid model-data driven approach."Energy Storage Materials 36 (2021): 435-458.

[10] Mazzi, Yahia, et al. "State of charge estimation of an electric vehicle's battery using tiny neural network embedded on small microcontroller units."International Journal of Energy Research 46.6 (2022): 8102-8119.

[11] Zhao, Kai, et al. "Digital twin-driven estimation of state of charge for Li-ion battery." 2022 IEEE 7th International Energy Conference (ENERGYCON). IEEE, 2022.

[12] Catelani, Marcantonio, et al. "Remaining useful life estimation for prognostics of lithium-ion batteries based on recurrent neural network."IEEE Transactions on Instrumentation and Measurement 70 (2021): 1-11.

[13] Halteren, Hans van, Jakub Zavrel, and Walter Daelemans. "Improving accuracy in word class tagging through the combination of machine learning systems."Computational linguistics 27.2 (2001): 199-229.

[14] Finn, Chelsea, Pieter Abbeel, and Sergey Levine. "Model-agnostic meta-learning for fast adaptation of deep networks."International conference on machine learning. PMLR, 2017.

[15] https://www.fmv.se/projekt/torped-47/

[16] Schmitt, Jakob, Ivo Horstkötter, and Bernard Bäker. "Effective estimation of battery state-of-health by virtual experiments via transfer-and meta-learning."Journal of Energy Storage 63 (2023): 106969.