Mean Variance Estimation Neural Network Particle Filter for Predicting Battery Remaining Useful Life

Francesco Cancelliere[1], Sylvain Girard [2], Jean-Marc Bourinet [3], Piero Baraldi [4], Enrico Zio [5]

[1,2] *PHIMECA, Paris, 75012, France*
cancelliere@phimeca.com
girard@phimeca.com

[1,3] *SIGMA Clermont University, Aubiere, 63178 , France*
jean-marc.bourinet@sigma-clermont.fr

[4,5] *Politecnico di Milano, Milano, 20161 , Italy*
piero.baraldi@polimi.it
enrico.zio@polimi.it

[5] *MINES Paris, PSL Research University, 06904 Sophia Antipolis, France*

ABSTRACT

Traditional remaining useful life (RUL) prediction methods based on particle filter (PF) require the manual tuning of hyperparameters, such as process or measurement noise, which poses challenges, particularly in real-life applications where external and operating conditions may change, potentially leading to large errors in the predictions. We address this issue by replacing the measurement equation of a PF with a mean variance estimation neural network that estimates the mean and the variance of the output distribution. As a result, the measurement noise is automatically estimated by the neural network and does not require manual setting. Through simulations and comparative analyses with state-of-the-art methods, the proposed mean variance estimation neural network particle filter (MVENN-PF) is shown to provide more stable and accurate RUL predictions, thereby potentially enhancing the robustness of battery health management systems based on it. Additionally, by eliminating the need to manually set a model hyperparameter (the measurement noise) the proposed method simplifies the modeling process, making it more accessible and adaptable to various battery systems.

1. INTRODUCTION

Predicting the remaining useful life of batteries allows optimizing maintenance schedules, and, therefore, reducing downtime and avoiding unexpected failures (Tran et al., 2022). This is a fundamental step for ensuring the longevity and optimal performance of batteries (Hu, Xu, Lin, & Pecht, 2020), which is a key element for the safety, cost efficiency, and sustainability of several industrial systems, such as electric vehicles and renewable energy storage (Chen et al., 2021).

Traditional model-based methods for RUL prediction rely heavily on physical models, which are not always available or accurately reflective of real-world conditions (Zio, 2022). Considering battery degradation, they may struggle to adapt to the stochastic nature of the process, which is influenced by numerous factors, including usage patterns, environmental conditions, and manufacturing variations. Particle filters (PF) (Kantas, Doucet, Singh, & Maciejowski, 2009) have been widely used for state estimation in dynamic systems. They suffer from issues such as particle degeneracy and the need of setting the hyperparameters process and measurement noise.

Data-driven approaches based on machine learning techniques use signal measurements to predict component present and future state of health and RUL (Wu, Fu, & Guan, 2016). They have been shown able to learn complex degradation patterns from historical data (Wang et al., 2021). Also, they are more adaptable and accurate in case of variation of external and operating conditions. However, they typically require large amounts of data to generalize effectively and struggle to describe unseen situations.

Hybrid methods such as physics-informed neural networks (PINNs) (Nascimento, Corbetta, Kulkarni, & Viana, 2021) and the combination of filter algorithms with machine learn-

ing approaches (Gu, Mahmoud, Arun, Liu, & Ma, 2023) have been proposed to partially address these problems. They integrate surrogate models with physical information, reducing the need for both expensive numerical models and large amounts of data. Specifically, data-driven particle filters combining a PF with a neural network (NN) offer promising perspective in terms of adaptability and accuracy of the predictions (Cadini, Sbarufatti, Cancelliere, & Giglio, 2019; Ma, Karkus, Hsu, & Lee, 2020). However, concerns remain on the robustness of these methods, whose performance are remarkably influenced by the values of some hyperparameters, difficult to manually set, especially for components working in variable operating and external conditions.

The method proposed in this paper addresses these limitations by enhancing a multi-layer perceptron particle filter (MLP-PF) (Cancelliere, Girard, & Bourinet, 2024). The main novelty is the replacement of the MLP with a mean variance estimation neural network (MVENN), which estimates both expected value and variance of the output (Nix & Weigend, 1994). Other works, such as (Nguyen, Medjaher, & Gogu, 2022), have utilized neural networks—specifically a long short-term memory (LSTM) network—to estimate the mean and variance of a lognormal distribution directly for RUL prediction. In our approach, the estimated variance is leveraged to model the measurement noise in the observation, rather than directly estimating the RUL. This strategy eliminates the need for manual tuning of measurement noise variance, which is one of the key hyperparameter of the MLP-PF and improving the robustness and accuracy of RUL predictions.

The key contributions of this paper include:

- Introducing an innovative data-driven particle filter (MVENN-PF) that by dynamically estimating the observation measurement noise eliminates the need of manually tuning this hyperparameter.
- Demonstrating the improved performance and robustness of the proposed method by comparing it with state-of-the-art methods on the NASA battery dataset (Saha & Goebel, 2007).

The proposed method is shown to improve the accuracy of RUL predictions for Li-Ion batteries and provide a robust framework adaptable to other prognostic applications. The paper is organized as follows: Section 2 provides a concise overview of data-drive particle filters, followed by the introduction of the novel MVENN-PF method. Section 3 describes the experimental data, introduces evaluation metrics, and presents the obtained results. Finally, Section 4 draws conclusions and discusses future applications of this method.

2. PROPOSED METHOD

Particle filters, also known as sequential Monte Carlo methods, are algorithms used to estimate the posterior probability density function (PDF) of a hidden state x_k at time step k, given a series of noisy observations $z_{0:k-1}$ (Doucet, Godsill, & Andrieu, 2000). The evolution of the state-space is described by the hidden Markov model:

$$x_k = f(x_{k-1}, \omega_{k-1}) \tag{1}$$

$$z_k = h(x_k, \eta_k) \tag{2}$$

where f and h represent the process and measurement equations, and ω and η denote process and measurement noise, respectively.

To recursively estimate the posterior PDF $p(x_k|z_{0:k})$, the PF follows a prediction-update cycle. First, a prior distribution $p(x_k|z_{0:k-1})$ is computed by generating N_s state trajectories, or particles, using the process equation. The distribution is then updated by assigning an importance weight to each particle i considering the likelihood $\mathcal{L} = p(z_k|x_k^i)$. Assuming a measurement equation with a Gaussian noise, the likelihood is:

$$\mathcal{L}_k^{(i)} = ((2\pi)^{k+1}|\Sigma_\eta|)^{-0.5} \exp\left\{-\frac{1}{2}\left(z_{0:k} - g(x_k^i, 0:k)\right)^T \Sigma_\eta^{-1}\left(z_{0:k} - g(x_k^i, 0:k)\right)\right\} \tag{3}$$

where Σ_η is the covariance matrix of the measurement noise.

Particle filters often face the degeneracy problem, where most particles after few prediction-update cycles have negligible weights. This can be mitigated by using ad-hoc techniques such as the sampling importance resampling (SIR), which resamples particles based on their weights to maintain diversity (Arulampalam, Maskell, Gordon, & Clapp, 2002). For a comprehensive overview of particle filters, the interested reader can refer to the seminal papers (Doucet et al., 2000; Arulampalam et al., 2002; Kantas et al., 2009).

2.1. Data-Driven Particle Filter

In situations where physical models are unavailable, a surrogate model can be employed within a PF (de Freitas, Niranjan, Gee, & Doucet, 2000). This approach adapts the state-space formulation to estimate the parameters x of a surrogate model $g(\cdot)$ that replaces the measurement equation. Consequently, the evolution of the state-space equations Eq. (1) and Eq. (2) in the data-driven framework is redefined as:

$$x_k = x_{k-1} + \omega_{k-1} \tag{4}$$

$$z_k = g(x_k, k) + \eta_k \tag{5}$$

The process model is a random perturbation of the parameters, controlled by the Gaussian process noise ω_{k-1}. In Eq. (5), η_k is the measurement noise associated to the measurement equation.

A possible implementation of data-driven PF uses a MLP as surrogate model. Indeed, the data-driven PF effectively es-

timates the MLP weights and biases when the underlying physical processes are not well understood, leveraging data to adaptively refine predictions. For example (Cancelliere, Girard, Bourinet, & Broggi, 2023) and (Jules, Cancelliere, Mattrand, & Bourinet, 2023) used simple MLP with one or two hidden layers, pre-trained on historical degradation trajectories, to predict the RUL of batteries and turbofans, respectively.

In these implementations, x_k contains the trained weights and biases of the NN, which are perturbed following Eq. (4) and resampled by applying Eq. (5) and Eq. (3). As a result, the particles with degradation trends similar to the observations are resampled while those far from the observations are discarded. This ensures that the particles converge to the observed degradation trend, adapting to changes in the external and operating conditions (Cadini, Sbarufatti, Cancelliere, & Giglio, 2019). A common problem of these applications is the selection of the proper process and measurement noises ω_k and η_k, which significantly influence performance. An approach to partially solve this issue is the use of a jittering variance (Jules et al., 2023) for the process noise ω_k in Eq. (4):

$$\omega_k = (\sigma_0 e^{-\frac{k}{\sigma_1}} + \sigma_2) \cdot ESS_k^{-\frac{1}{3}} \tag{6}$$

where $\sigma_0 e^{-\frac{k}{\sigma_1}} + \sigma_2$ is a decreasing variance, commonly used in PF (Sbarufatti, Corbetta, Giglio, & Cadini, 2018; Cadini, Sbarufatti, Corbetta, Cancelliere, & Giglio, 2019), and $\sigma = [\sigma_0, \sigma_1, \sigma_2]$ is a vector of hyperparameters to be tuned. However, the importance of tuning σ is decreased by the introduction of the effective size sample (ESS) term (Flury & Shephard, 2009):

$$ESS = \frac{1}{\sum_{i=1}^{N_s} (\hat{w}_i^k)^2} \tag{7}$$

ESS is an evaluation of the degeneracy of the particles. When all particles have similar weights, ESS approaches N_s, whereas if most particles have weights close to zero and only a few have significant weights, ESS will be small. According to Eq. (6), a small ESS increases the jittering variance, leading to a larger exploration of the state space, while a large ESS decreases this capability of exploration.

In (Sbarufatti et al., 2018; Cancelliere et al., 2023), the measurement noise η_k is assumed zero-mean, Gaussian with variance σ_η^2. The tuning of σ_η is based on a "rule of thumb", and it is assumed independent from k. However, the impact of this term on the performance of the algorithm is remarkable, as it is used to compute the likelihood in Eq. (3). For example, Fig. 1 shows the result of three simulations which used the same data and different values of σ_η. It can be observed that when a small σ_η value is employed (orange lines) the particles are not able to effectively follow the degradation trend. This is primarily due to the inability of the likelihood function to differentiate between accurate and inaccurate predictions, resulting in similar weights assigned to all particles. Consequently, this scenario resembles a random perturbation of the NN parameters x_k, leading to unreasonable predictions of the degradation evolution. When σ_η is large (green lines), it leads to broad likelihood distributions. This implies that large likelihoods are associated to many particles. Consequently, the filter predicts large confidence intervals, which can confuse the maintenance decision-maker. Finding a good intermediate value of σ_η can be a case-dependent challenging task. Also, in case of modification of the operating and environmental conditions, a new setting of the parameter may be required.

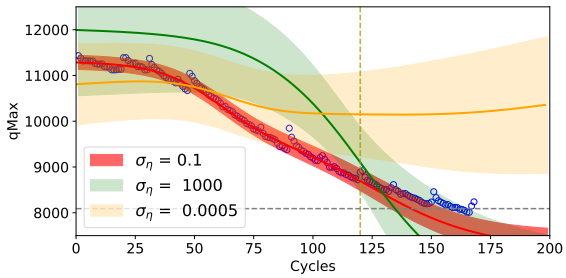

Figure 1. Predictions of the evolution of the component degradation (q_{Max} (Cancelliere et al., 2023)) performed at cycle 120 by MLP-PF with different measurement noises. The blue circles represent the ground truth degradation values, the shaded areas correspond to the 95th percentile confidence intervals of the predictions, while the solid lines represent the prediction expected value.

2.2. Neural Network for variance estimation

Aiming to increase the robustness of the algorithm and avoid the issue of tuning σ_η, we propose a NN architecture that directly estimates the variance of its output, and, therefore, allows automatically estimating the measurement noise. The architecture is based on the work by (Nix & Weigend, 1994), which introduced a method to estimate both the mean and the variance of a target probability distribution. Fig. 2 shows the traditional NN architecture and the architecture in (Nix & Weigend, 1994), where a new auxiliary output unit is added to estimate the uncertainty of the prediction. Specifically, the added output neuron estimates the variance of the output PDF. Considering a set of N_t input-output training patterns $(k, z_k)_{k=1,.,N_t}$, the loss function used to train the NN is the negative log-likelihood:

$$\text{NLL} = \frac{1}{2}\left(\log(\hat{s}^2(k)) + \frac{(z_k - \hat{g}(k))^2}{\hat{s}^2(k)}\right) \tag{8}$$

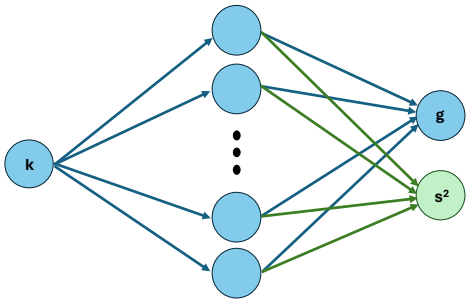

Figure 2. Comparison of the neural network architectures of the MLP-PF (Cancelliere et al., 2023) (blue) and of the proposed MVENN-PF with the added output neuron (green).

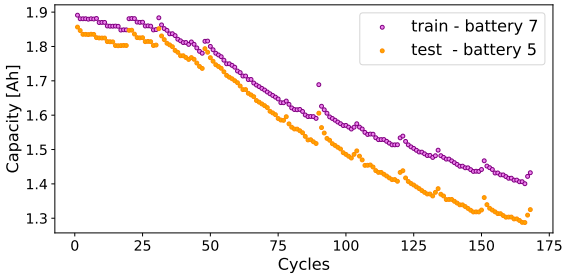

Figure 3. Evolution of the capacity during the lifetime of two batteries. The data collected from battery 7 are used for model training, whereas the data from battery 5 are used for testing the algorithm.

where $\hat{g}(k)$ is the prediction of the expected value of the output and $\hat{s}^2(k)$ is the corresponding estimated variance. It is interesting to note that by embedding this NN in the MLP-PF the method allows estimating a measurement noise of variance $\hat{s}^2(k)$ that can change with time. This is because the NN provides as output a value of $\hat{s}^2$ for each input k.

The integration of the proposed neural network architecture into the data-driven PF requires to address two practical issues. Firstly, the shallow architecture with 3 neurons in a single hidden layer suggested in (Cancelliere et al., 2023) for the data-driven PF has been shown not able to capture the degradation dynamics and the associated noise. For this reason, in this study we opted for a more complex architecture formed by a single layer with 20 hidden neurons. The increased complexity of the NN architecture translates into a larger number of weights and biases to be packed in x_k and to be recursively estimated as new observations become available. To counterbalance the larger number of parameters, the variance of the noise ω_k of the process equation (Eq. (4)) is reduced to maintain stability. Essentially, when the number of model parameters increases, the same process noise ω causes larger variations of weights and biases of the NN output, which can lead to difficulty in the convergence of the PF algorithm.

Lastly, in the computation of the likelihood $\mathcal{L}$, a notable difference emerges with respect to the covariance matrix Σ_η. Differently from the approach in (Cancelliere et al., 2023), where the covariance matrix Σ_η contains the same value ω_η on the main diagonal (the same noise is applied to all particles at each time k), the proposed formulation uses a different matrix Σ_η^i for each i^{th} particle, with the k^{th} term on the main diagonal being $\sigma_{\eta,k}^i$, equal to the NN output value $s^2(x_k^i)$ where x_k^i is the state of particle i at time step k.

3. EXPERIMENTAL SETUP AND RESULT

To test the robustness and validity of the proposed method we apply it to the well-known NASA battery dataset (Saha & Goebel, 2007). The dataset contains data about the degradation of batteries during their life cycles. We consider the capacity measures, which are computed at each battery discharge cycle, and represent the amount of electricity (A/h) the battery is able to provide when a load is applied, and can be considered as a degradation indicator. In particular, we use battery 7 data for the initial training of the network and we use battery 5 data to evaluate the performance of the proposed method in terms of RUL prediction (Fig. 3). It can be observed that the two batteries, despite being subjected to identical and controlled laboratory conditions (including constant load and external temperature), show different degradation patterns. Indeed, the difference in the remaining capacity after 175 cycles is significant, with battery 7 which maintained a capacity over 1.4 Ah, whereas battery 5 capacity dropped below 1.3 Ah. This difference in end-of-life capacities underscores the variable nature of Li-Ion battery degradation, even under uniform external conditions, and highlights the importance of developing prognostic methods able to promptly adapt to the incoming real-time observations. Unlike traditional neural network approaches that require large datasets for effective training, in this work, we used a single degradation trajectory to train the neural network, as it serves merely for initialization purposes. This significantly reduces data requirements and is advantageous compared to conventional methods where performance is highly dependent on data volume.

The proposed MLP network architecture, illustrated in Fig. 2, takes as input the time step value k, and has 2 output neurons: the corresponding capacity and the estimated measurement noise variance s^2. The hidden layer is composed of 20 neurons, using tanh as activation function. This architecture consists of 82 weights and biases to be recursively estimated by the particle filter. The three variances of the process noise in Eq. (6) have been selected: $\sigma_0 = 5 \times 10^{-5}$, $\sigma_1 = 10^2$ and $\sigma_2 = 10^{-6}$.

As comparison method we use the same network architecture and hyperparameters of (Cancelliere et al., 2023), which consists of a MLP with a single hidden layer of 3 nodes, resulting

Table 1. Prognostic Evaluation Metrics.

	MLP-PF	MVENN-PF
β	0.388	**0.578**
β 25	0.236	**0.286**
CRA	0.667	**0.758**
CRA 25	0.419	**0.560**
CIC	0.494	**0.650**
CIC 25	0.780	**0.951**

in 10 weights and biases. The variances of the process noise (Eq. (6)) have been set equal to $\sigma_0 = 5 \times 10^{-3}$, $\sigma_1 = 10^2$ and $\sigma_2 = 10^{-4}$, which are considerably larger than those used in the proposed method. The measurement noise variance has been assumed $\sigma_\eta = 10^{-1}$. For consistency, for both cases we used the same number of particles $N_s = 500$.

3.1. Prognostic Metrics

The performances are evaluated using the α-λ plot and the β metric (Lall, Lowe, & Goebel, 2010), the cumulative relative accuracy (CRA) (Saxena et al., 2008) and the confidence interval coverage (CIC) (Jules et al., 2023).

The α-λ plot compares the accuracy of the prediction with respect to the ground truth remaining useful life (RUL^{GT}). λ represents the normalized time and is defined as $\lambda = (k/T_{fail})$, where k is the time step and T_{fail} is the time step associated to the end of life. Hence, $\lambda = 1$ corresponds to the end of life of the battery. The α-bounds, calculated at each time step k as $RUL \pm \alpha$, represent a goal region for the prediction to be considered as successful. In this work we use $\alpha = 0.2$. The β metric is computed starting from the α-λ plot, and is defined as the area under the predicted RUL pdf that falls within the α bounds:

$$\beta = \frac{1}{T_{fail}} \sum_{k=0}^{T_{fail}} \int_{\text{RUL}_k - \alpha}^{\text{RUL}_k + \alpha} \text{PDF}(\widehat{\text{RUL}})\, d\text{RUL} \tag{9}$$

This metric allows to discriminate between predictions with different levels of associated uncertainty: a high β metric indicates a superior RUL prediction.

The CRA represents the distance of the prediction to the ground truth EOL, evaluated at each time step. It is defined as:

$$\text{CRA} = \frac{1}{T_{fail}} \sum_{k=0}^{T_{fail}} \left(1 - \left| \frac{\text{RUL}_k^{\text{GT}} - \text{RUL}_k^{\text{pred}}}{\text{RUL}_k^{\text{GT}}} \right| \right) \tag{10}$$

A perfect prediction has a value of 1. This metric emphasizes errors closer to the failure of the battery, highlighting the importance of making good predictions close to the EOL.

Finally, the CIC is used to assess the prediction considering the confidence interval:

$$CIC = \frac{1}{T_{fail}} \sum_{k=0}^{T_{fail}} \mathbb{1}_{\text{RUL}^{\text{GT}} \in \widehat{CI_k}} \tag{11}$$

where $\mathbb{1}_{RUL_k^{GT} \in \widehat{CI_k}}$ is the indicator function, which is equal to 1 if the ground truth RUL lies within the predicted confidence interval and 0 otherwise. If the prediction interval includes the RUL^{GT} at each time step, the CIC is 1, while is 0 if the ground truth RUL is always outside the confidence interval.

3.2. Results

Fig. 4 compares the predictions of the capacity evolution provided by the proposed method (MVENN-PF) with those of the MLP-PF in (Cancelliere et al., 2023), performed at different time steps. As expected, the initial prediction of battery 5 capacity provided by the two methods are very similar, and characterized by a capacity significant overestimation, since both NNs are trained using the same data taken from battery 7, which has a longer lifetime than the battery 5. Nonetheless, already at cycle 50 ($\lambda = 0.31$) both methods adapt the predictions to the incoming measurements and are making more accurate predictions. Fig. 5 shows the RUL predictions provided by the two methods. Note that MLP-PF updates its predictions earlier than the proposed method due to the larger initial value of ω_k. In contrast, the MVENN-PF algorithm exhibits more stable behavior, adapting more cautiously due to its capability of dynamically estimating the measurement noise. The average variance of the measurement noise estimated by the NN is larger than the hyperparameter σ_η used by MLP-PF, indicating that the proposed method anticipates noisier observations and therefore adjusts more gently the predictions.

This behavior is more evident when comparing the simulation between cycle 50 (Fig. 4b) and cycle 80 (Fig. 4c). During this phase, battery 5 shows a pronounced degradation trajectory. The MLP-PF algorithm is overly reactive to this trend, resulting in an underestimated prediction of the EOL. The MVENN-PF algorithm, on the other hand, remains stable, providing a consistent prediction and accurately estimating the RUL.

The metrics, summarized in Table 1, confirm the superior prognostic performance of the proposed method. These metrics are calculated over the entire battery life and specifically for the last quarter, which is the most critical phase for prognostics. The results demonstrate that the proposed method outperforms MLP-PF, providing more accurate and stable predictions of the RUL. This improved performance is crucial for effective battery management and underscores the reliability of the new method.

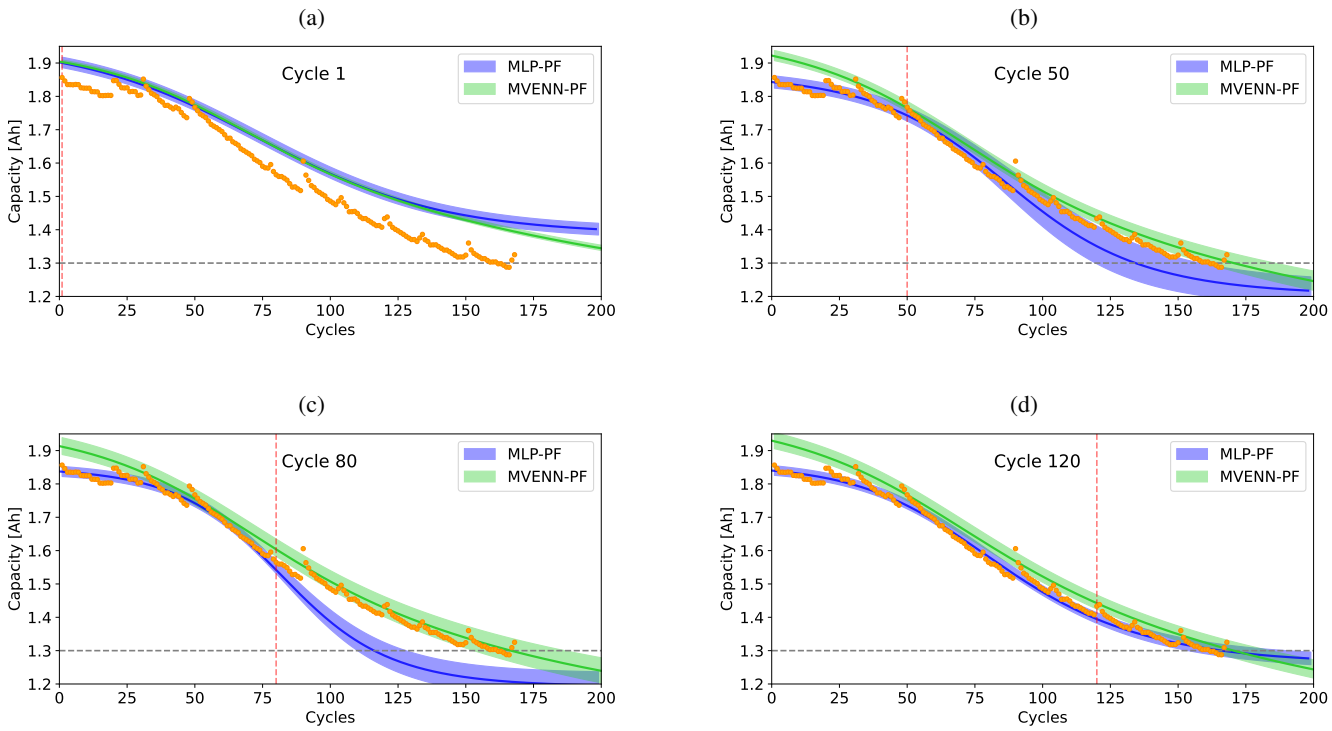

Figure 4. Predictions of the evolution of battery 5 capacity performed at cycles 1 (a), 50 (b), 80 (c) and 120 (d) by the proposed method (MVENN-PF) and the comparison method (MLP-PF). The shaded areas correspond to the 95th percentile confidence intervals of the predictions, while the solid lines represent the prediction expected value.

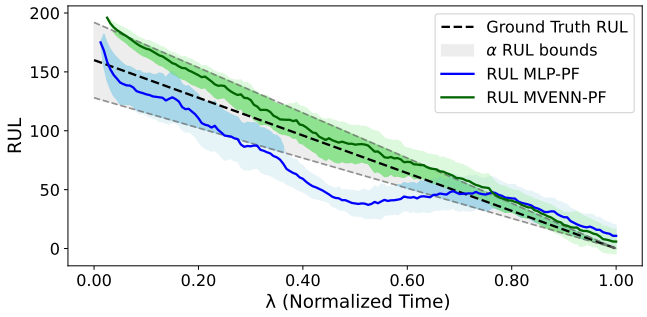

Figure 5. Comparison of results in terms of RUL

4. CONCLUSION

In this paper, we have presented a novel method for enhancing the prediction of battery remaining useful life by combining a data-driven particle filter with a neural network with variance estimation (MVENN-PF). The experimental results indicate that the MVENN-PF algorithm outperforms a traditional MLP-PF in terms of key prognostic metrics, including confidence interval coverage (CIC), cumulative relative accuracy (CRA), and β metric.

The proposed method offers two main contributions:

- **Elimination of σ_η hyperparameter tuning**: One of the significant advantages of the MVENN-PF method is that it removes the need to manually tune the variance of the measurement noise σ_η, which greatly affects the performance of the classical algorithms. This self-adjusting feature simplifies the implementation process and enhances the algorithm adaptability to different datasets and operational conditions.
- **Improved robustness to measurement noise**: The inherent robustness of the MVENN-PF algorithm is another key contribution. By providing an estimate of the measurement noise variance through the NN, the algorithm dynamically adjusts to varying noise levels, leading to more stable and reliable RUL predictions. This capability ensures that the algorithm maintains high performance even in the presence of noisy data, which is crucial for real-world applications.

These advancements enhance the reliability of particle filter-based prognostic algorithms, which often suffer from robustness issues, and contribute to the development of more robust and adaptive health management solutions for battery-operated systems. Moreover, the real-time capability of this algorithm is ensured by its computational efficiency: with the proposed MVENN-PF architecture, the execution time per

cycle on a standard commercial laptop never exceeds 0.1 seconds.

Future work will focus on further refining the NN architecture and exploring the integration of additional data sources to enhance the robustness and accuracy of RUL predictions. Additionally, the robustness of the algorithm will be tested in scenarios involving biased or broken sensors to verify if the estimation of the measurement noise can cope with these extreme situations.

ACKNOWLEDGMENT

The project leading to this application has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 955393

The work of Piero Baraldi is supported by FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence).

REFERENCES

Arulampalam, M. S., Maskell, S., Gordon, N., & Clapp, T. (2002). A tutorial on particle filters for online nonlinear/nongaussian bayesian tracking. *IEEE Transaction on Signal Processing*, *50*(2), 174–188. doi: 10.1109/9780470544198.ch73

Cadini, F., Sbarufatti, C., Cancelliere, F., & Giglio, M. (2019). State-of-life prognosis and diagnosis of lithium-ion batteries by data-driven particle filters. *Applied Energy*, *235*(June 2018), 661–672. doi: 10.1016/j.apenergy.2018.10.095

Cadini, F., Sbarufatti, C., Corbetta, M., Cancelliere, F., & Giglio, M. (2019). Particle filtering-based adaptive training of neural networks for real-time structural damage diagnosis and prognosis. *Structural Control and Health Monitoring*, *26*(12), 1–19. doi: 10.1002/stc.2451

Cancelliere, F., Girard, S., & Bourinet, J.-M. (2024). Data-Driven Prognostics with Multi-Layer Perceptron Particle Filter: a Cross-Industry Exploration. *PHM Society European Conference*, *8*(1), 8. doi: 10.36001/phme.2024.v8i1.4034

Cancelliere, F., Girard, S., Bourinet, J.-M., & Broggi, M. (2023). Grey-box Approach for the Prognostic and Health Management of Lithium-Ion Batteries. *Annual Conference of the PHM Society*, *15*(1), 1–8. doi: 10.36001/phmconf.2023.v15i1.3506

Chen, Y., Kang, Y., Zhao, Y., Wang, L., Liu, J., Li, Y., ... Li, B. (2021). A review of lithium-ion battery safety concerns: The issues, strategies, and testing standards. *Journal of Energy Chemistry*, *59*, 83–99. doi: 10.1016/j.jechem.2020.10.017

de Freitas, J. F., Niranjan, M., Gee, A. H., & Doucet, A. (2000). Sequential Monte Carlo methods to train neural network models. *Neural computation*, *12*, 955–993.

Doucet, A., Godsill, S., & Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, *10*, 197–208. doi: 10.1023/A:1008935410038

Flury, T., & Shephard, N. (2009). Learning and filtering via simulation : smoothly jittered particle filters. *RECALL*(2001), 1–27.

Gu, H., Mahmoud, H., Arun, R. L., Liu, J., & Ma, X. (2023). Particle Filter and Its Variants for Degradation State Estimation and Remaining Useful Life Prediction. *Proceedings - 2023 Prognostics and Health Management Conference - Paris, PHM-Paris 2023*, 256–263. doi: 10.1109/PHM58589.2023.00055

Hu, X., Xu, L., Lin, X., & Pecht, M. (2020). Battery Lifetime Prognostics. *Joule*, *4*(2), 310–346. doi: 10.1016/j.joule.2019.11.018

Jules, E., Cancelliere, F., Mattrand, C., & Bourinet, J.-M. (2023). Remaining useful life prediction of turbofans with virtual health indicator: A comparison of particle filter-based approaches. , 75-82. doi: 10.1109/ICSRS59833.2023.10381439

Kantas, N., Doucet, A., Singh, S., & Maciejowski, J. (2009). An Overview of Sequential Monte Carlo Methods for Parameter Estimation in General State-Space Models. In *Ifac symposium on system identification* (pp. 774–785). IFAC. doi: 10.3182/20090706-3-FR-2004.00129

Lall, P., Lowe, R., & Goebel, K. (2010). Prognostics using Kalman-Filter models and metrics for risk assessment in BGAs under shock and vibration loads. *Proceedings - Electronic Components and Technology Conference*(1), 889–901. doi: 10.1109/ECTC.2010.5490691

Ma, X., Karkus, P., Hsu, D., & Lee, W. S. (2020). Particle filter recurrent neural networks. *AAAI 2020 - 34th AAAI Conference on Artificial Intelligence*, 5101–5108. doi: 10.1609/aaai.v34i04.5952

Nascimento, R. G., Corbetta, M., Kulkarni, C. S., & Viana, F. A. (2021). Hybrid physics-informed neural networks for lithium-ion battery modeling and prognosis. *Journal of Power Sources*, *513*(August), 230526. doi: 10.1016/j.jpowsour.2021.230526

Nguyen, K. T., Medjaher, K., & Gogu, C. (2022). Probabilistic deep learning methodology for uncertainty quantification of remaining useful lifetime of multi-component systems. *Reliability Engineering and System Safety*, *222*(February), 108383. Retrieved from https://doi.org/10.1016/j.ress.2022.108383 doi: 10.1016/j.ress.2022.108383

Nix, D. A., & Weigend, A. S. (1994). Estimating the Mean and Variance of the Target Probability Distribution. In

Proceedings of 1994 ieee international conference on neural networks (icnn'94) (pp. 55–60). Orlando, FL. doi: 10.1109/ICNN.1994.374138

Saha, B., & Goebel, K. (2007). *Battery Data Set. NASA Ames Prognostics Data Repository.*

Saxena, A., Celaya, J., Balaban, E., Goebel, K., Saha, B., Saha, S., & Schwabacher, M. (2008). Metrics for evaluating performance of prognostic techniques. *2008 International Conference on Prognostics and Health Management, PHM 2008*(October). doi: 10.1109/PHM.2008.4711436

Sbarufatti, C., Corbetta, M., Giglio, M., & Cadini, F. (2018). Adaptive prognosis of lithium-ion batteries based on the combination of particle filters and radial basis function neural networks. *Journal of Power Sources*, *344*, 128–140. doi: 10.1016/j.jpowsour.2017.01.105

Tran, M. K., Mevawalla, A., Aziz, A., Panchal, S., Xie, Y., & Fowler, M. (2022). A Review of Lithium-Ion Battery Thermal Runaway Modeling and Diagnosis Approaches. *Processes*, *10*(6). doi: 10.3390/pr10061192

Wang, S., Jin, S., Bai, D., Fan, Y., Shi, H., & Fernandez, C. (2021). A critical review of improved deep learning methods for the remaining useful life prediction of lithium-ion batteries. *Energy Reports*, *7*, 5562–5574. doi: 10.1016/j.egyr.2021.08.182

Wu, L., Fu, X., & Guan, Y. (2016). Review of the Remaining Useful Life Prognostics of Vehicle Lithium-Ion Batteries Using Data-Driven Methodologies. *Applied Sciences*, *6*(6), 166. doi: 10.3390/app6060166

Zio, E. (2022). Prognostics and Health Management (PHM): Where are we and where do we (need to) go in theory and practice. *Reliability Engineering and System Safety*, *218*(PA), 108119. doi: 10.1016/j.ress.2021.108119