STATE-OF-LIFE PROGNOSIS AND DIAGNOSIS OF LITHIUM-ION BATTERIES BY DATA-DRIVEN PARTICLE FILTERS

F. Cadini[1], C. Sbarufatti[1], F. Cancelliere[2], M. Giglio[1]

[1]Dipartimento di Meccanica, Politecnico di Milano, Italy

[2]Dipartimento di Energia, Politecnico di Milano, Italy

francesco.cadini@polimi.it

Abstract

The aim of this study is that of presenting a new diagnostic and prognostic method aimed at automatically detecting deviations from the expected degradation dynamics of the batteries due to changes in the operating conditions, or, possibly, anomalous behaviors, and predicting their remaining useful life (RUL) in terms of their state-of-life (SOL), without needing to derive any complex physics-based models and/or gather huge amounts of experimental data to cover all possible operative/fault conditions. The proposed method in fact exploits the real time framework offered by particle filtering and resorts to neural networks in order to build a suitable parametric measurement equation, which provides the algorithm with the capability of automatically adjusting to different battery's dynamic behaviors. The results of this study demonstrate the satisfactory performances of the algorithm in terms of adaptability and diagnostic sensibility, with reference to suitably identified case studies based on actual Lithium-Ion battery capacity data taken from the prognostics data repository of the NASA Ames Research Center database and of the CALCE Battery Group.

1. INTRODUCTION

In recent years, Lithium-ion (Li-ion) batteries have gained large popularity as portable energy sources due to their significant advantages with respect to other battery types, such as: i) the lower weight, due to the lightweight lithium and carbon-made electrodes, and, at the same time, the larger energy density, due to the high chemical reactivity of lithium; ii) the possibility of being recharged also if they are not completely discharged without any detrimental effects (no memory effect); iii) the lower

self-discharge rate, so that they better maintain their charge when not used; iv) the longer life cycle, since they can operate successfully even after hundreds of charge-discharge cycles [1], [2]. These features have played a major role in their successful use in many different application fields, including consumer electronics (cell phones, laptops, etc.) [ref b], hybrid and electric vehicles in the automotive industry [ref g], the more recent development of hybrid/electric aircraft (airplanes and helicopters), balanced management of electric power grids with significant contributions from fluctuating renewable power sources (solar and wind) [ref d] and [ref e], where studies are carried on aiming even at exploiting, at a domestic scale, Li-ion batteries previously used in electrical vehicles and no more satisfying the requirements of the automotive industry [ref c] and space exploration manned and unmanned missions [ref h].

However, due to their rechargeable nature, Li-Ion batteries are subject to irreversible processes occurring during their charging and discharging cycles, such as, for example, the formation of a solid-electrolyte interphase (SEI) [3], which severely affect the batteries' electrochemistry. These processes involve, in general, a continuous capacity fade, which eventually leads to the battery failure, with consequences ranging from a quite safe need to replace the battery of a mobile phone, to the catastrophic failure of an interplanetary probe [1], [4].

In order to overcome these issues, many efforts have been devoted in literature to devising proper methods for improving the reliability and the availability of Li-Ion batteries. More specifically, a major role is played by the so-called prognostic and health management (PHM) methods, which, on the basis of different kinds of available, but indirect, information, allow to automatically, and in real-time, track some hidden indicators of the degradation state of the batteries, such as, for example the state-of-health (SOH), the state-of-charge (SOC), the state-of-life (SOL), and at predicting their remaining useful lifetime (RUL), either in terms of the end-of-discharge (EOD) or the end-of-life (EOL) times, possibly to support condition- or even prediction-based maintenance policies. In this regard, thorough reviews of many advanced PHM methods can be found in [5], [6]. Traditionally, these methods are classified in three major families, i.e., model (or physics)-based, data-driven or hybrid methods, depending on the type and quality of the information used to perform diagnostics and/or prognostics [7]. Model-based methods focus on identifying proper relationships between the observable quantities and the indicators of interest by building physics-based models of the degradation processes affecting the battery life [8]–[14]. Data-driven methods, on the other hand, aim at mapping the above by some approximating models adaptively built on the basis of available data, such as, for example, neural networks (NN), Gaussian process functional regressions, support vector regressions, fuzzy inference engines, etc [15]–[19]. Hybrid methods aim at combining model-based and data-driven methods, when possible, in an attempt to overcome the limitations of the individual

methods and, thus, improve diagnostic and prognostic accuracies by better exploiting all the available information [20]–[24]. A promising hybrid strategy seems that of resorting to particle filtering-based algorithms, where the required analytical models representing either the dynamic behavior of the system or the measurement equation are actually suitable data-driven surrogate models [25], [26]. These kind of methods are based on the consideration that, both physics-based model and surrogate models require the identification of suitable model parameters on the basis of some available observations; however, surrogate models do not require any physics/analytic-based derivations, which might turn out to be very time consuming, and are generally much computationally faster, especially with respect to numerical models, which might be a critical feature for real-time applications (as one can devise by reading the thorough review of the state-of-the-art available physical models in [ref i]).

One important issue which still severely limits the applicability of these approaches is that the surrogate models are trained off-line on the basis of a set of available examples of the input/output mapping of interest, typically collected under some representative, fixed operative conditions, such as those offered by the controlled environments of laboratory tests. Actually, this represents a problem also for many other diagnostic/prognostic methods, not being restricted to those relying on surrogate modeling, although the latter mostly suffer from this limitation, due to the fact that their generalization capability only depends on the available data, and not on physical reasoning. For example, many works of literature derive and demonstrate their proposed methods with reference to sets of Li-ion batteries voltage/capacity laboratory measurements acquired at different, but constant discharge rates/currents and under controlled laboratory conditions (e.g. fixed ambient temperature, charge/discharge scheduled procedures, etc.). As the size of the available dataset increases, possibly including information on different operating conditions, this approach allows to devise algorithms with increasing prediction/generalization capabilities; at the same time, however, it does not account for operating conditions varying in real time, as typically occurring when considering the actual batteries mission profiles and/or boundary conditions (e.g., temperature, mechanical degradation, lithium metal plating and other ageing mechanisms, as summarized in [27]), thus severely limiting its application to real-life problems.

Some works of literature have already attempted to address this issue, which requires the capability of adapting in real time to the changing underlying dynamics. For example, in [28] the authors propounded an adaptive particle filter-based algorithm for predicting the remaining useful life of a battery when operating at different discharge rates; similarly, in [ref l], an equivalent circuit model accounting for the discharge current is used to perform adaptive diagnosis of Li-ion batteries under real operative mission profiles. Yet, the parameters of the resulting equivalent circuit models are still

identified off-line on the basis of available data (e.g., from AC Impedance Spectroscopy) and many of the factors actually influencing the degradation dynamics in real operative conditions are still not taken into account. In [2] some of the same authors of the present work presented a novel hybrid prognostics framework for the prediction of the end of discharge of Li-ion batteries, where the parameters of the surrogate model, i.e. a radial basis function neural network, were identified on-line by a particle filter on the basis of the real-time observations of the degradation process. That work represented a first attempt to create a prognostic tool capable of automatically coping to varying boundary conditions, with no need for explicitly modeling the dependency of the degradation dynamic on the external influencing factors. This allowed to naturally capture possible changes of the degradation dynamics and to accordingly update the RUL estimates. However, the algorithm was still only tailored to the specific EOD prognosis problem, and was still lacking of general optimization strategies which would guarantee its applicability also to different prognostic problems. Furthermore, the algorithm did not allow to perform any diagnostic tasks, which are fundamental for PHM applications.

In this context, the first purpose of this work is that of adapting the hybrid approach introduced in [2], which was restricted to the EOD prediction within individual discharge cycles, to be able to perform also SOL estimation and predict the EOL of Li-ion batteries. First, it is proposed to resort to multi layer perceptron (MLP) neural networks, which have turned out to be simpler and more intuitive for this kind of application. Similarly to [2], the parameters of the MLP networks are adaptively identified on-line by the particle filter on the basis of real-time observations of the Li-ion battery capacity. Then, since, in this case, the algorithm has to predict an individual degradation history, and not several successive discharge curves as in the previous work, a pre-training of the MLP neural network on some reference trajectory is suggested (although not strictly required), so as to significantly restrict the search space of the surrogate model's parameters and speed-up the algorithm convergence. Note that the actual degradation can be significantly different from the pre-training one, as it is demonstrated in this work, still guaranteeing the adaptation capabilities of the prognostic tool. Moreover, to further increase the algorithm adaptability, the set of particles used by the particle filter (i.e., neural networks weights, see [2]) is artificially enriched by a particle resulting from a back-propagation-based optimization of a network on the basis of the capacity observations available up to the current time.

The second objective of the present work is that of providing the proposed algorithm with an additional, original diagnostic module, based on the particle filter-based estimation of the observations log-likelihood ratio, for the automatic and on-line detection of any changes in the "expected" dynamic degradation process. This represents a very important task for Li-ion batteries,

which, to the authors' knowledge, appears to be somehow overlooked by Li-ion PHM specialists, but may significantly improve the prognostic capabilities and, more generally, the safe management of the batteries, especially when operating in real environments subject to varying boundary conditions. The proposed method is demonstrated with reference to real degradation transients taken from the NASA Ames Research Center database [29] and the CALCE Battery Group [30]. The available trajectories are also properly modified in order to be able to test the algorithm in varying operating conditions and to prove the effectiveness of its diagnostic module.

The paper is organized as follow. Section 2 briefly recalls the main features of the method proposed for sequentially train MLP-NN models by means of a particle filter algorithm. The multi layer perceptron-based particle filter (MLP-PF) approach here proposed is then customized in order to perform adaptive prognosis of the EOL of Li-Ion batteries and diagnosis of their SOH. Section 3 discusses the performances of the proposed method, demonstrating the capability of the algorithm with reference to the datasets cited above, which are typically used as benchmark case studies in similar works of literature. Section 4 draws some conclusions on the results and proposes future developments of the methodology.

2. MULTI LAYER PERCEPTRON PARTICLE FILTER (MLP-PF) FOR DIAGNOSIS AND PROGNOSIS

In this Section first the method originally proposed in [31] and further developed and applied to PHM of Li-ion batteries in [2] is shortly recalled, where particle filters were used to sequentially, on line, train neural network models. The multi layer perceptron-based particle filter (MLP-PF) approach here proposed is tailored to the problem under investigation in order to be able to predict the EOL and to diagnose the SOH of Li-Ion batteries subject to repeated charge-discharge cycles. The interested reader is referred to [32], and to [33] and [34] for thorough descriptions of the functioning of MLPs and particle filters, respectively, and to [31] for further details on particle filtering-based NN training.

2.1 Multi Layer Perceptron neural networks

Figure 1 shows the structure of the MLP networks adopted in this work, which are modeled by means of the NETLAB MATLAB package [35]. The MLP model aims at representing the degradation behavior (in terms of capacity) of the Li-ion batteries as a function of the number of charge-discharge cycles. The input node is fed by the number of cycles (k) at which the capacity is observed, whereas

the capacity observation (z) is associated to the output node, which, in turn, gathers i) the outputs from a generic number M of non-linear hidden nodes, each weighted by a factor $\theta_i^{(1,2)}$, and ii) the biases $b_i^{(1,2)}$, $i = 1, \dots, M$.

The network parameters $\theta_i^{(1,2)}$ and $b_i^{(1,2)}$, $i = 1, \dots, M$, are collected in the vectors $\boldsymbol{\theta} \subseteq \boldsymbol{\Theta} \in R^{n_\theta x 1}$ and $\boldsymbol{b} \subseteq \boldsymbol{B} \in R^{n_b \times 1}$, respectively. According to the classical MLP structure, the non-linear activation function of the hidden neurons $h(\cdot)$: $R^{1\,x\,1}_{[-\infty,\infty]} \rightarrow R^{1\,x\,1}_{[-1,1]}$ is a tan-sigmoid:

$$h\left(\theta_i^{(1)}\right) = \frac{2}{1 + e^{-2\theta_i^{(1)}}} - 1; \quad i = 1, \dots, M \qquad (1)$$

while the output activation function $f(\beta) = \beta$ (with β being the generic node input) is linear. Then, the output $g\big((\theta, b), k\big)$ of the MLP network represented in Figure 1 is:

$$g\big((\theta, b), k\big) = f\left(\sum_{i=1}^{M}\left(h\left(\left(k\theta_i^{(1)} + b_i^{(1)}\right)\theta_i^{(2)}\right) + b^{(2)}\right)\right) \qquad (2)$$

Note that, given a set of values for the parameters $\theta_i^{(1,2)}$ and $b_i^{(1,2)}$, $i = 1, \dots, M$, the MLP output $g\left(k, \theta_i^{(1,2)}, b_i^{(1,2)}\right)$ is an individual battery capacity history as a function of the discrete charge-recharge cycle index k.

A good choice of the number of hidden nodes M, should, intuitively, rely on the general consideration that the larger the number of hidden neurons is, the higher the capability of the algorithm in approximating complex input/output mappings [32]. On the other hand, too large a number of hidden neurons may result in data over-fitting issues, which severely hamper the generalization capability of the NN models. In addition to that, a large number of hidden neurons also implies a larger number of unknown parameters network weights which must be determined, which may significantly threaten the efficiency of the training/updating process described later.

Similarly to what was done in [2] for the RBF networks, in this application the number of hidden neurons is empirically chosen equal to 3 on the basis of a trial and error procedure. Similarly to [2], a qualitative sensitivity analysis performed by the authors (not shown here for brevity's sake) has shown that increasing the number of hidden nodes to 4-7 does not significantly modify the algorithm performances.

The $n_x = 10$ network parameters are then organized in a vector $\boldsymbol{x} = [x_1, x_2, \dots, x_{10}]^T$, where the elements $\boldsymbol{x}_{1:6}$ are the six connection weights ($\boldsymbol{\theta}_{1:6}$) and the elements $\boldsymbol{x}_{7:10}$ are the four MLP biases ($\boldsymbol{b}_{1:5}$). The MLP inputs k and the Li-ion battery capacity observations z_k are normalized with respect to the means and the standard deviations of the available data.

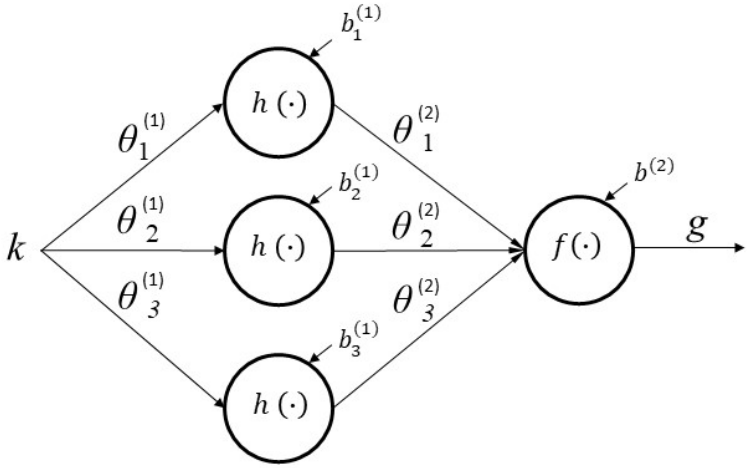

Figure 1: MLP neural network structure.

2.2 MLP-PF for Li-ion batteries EOL prognosis

Following the strategy proposed in [2], a state-space representation is here defined for the n_x MLP model parameters (weights and biases), which are then treated as unknown parameters to be sequentially estimated within the Bayesian framework offered by particle filtering. The MLP parameters are thus collected in the hidden state vector $\boldsymbol{x} \in \mathbb{R}^{n_x \times 1}$. Under the common operative assumption that the process observations z_k are available at discrete cycles k, the state space representation then can be written as:

$$\begin{aligned} \boldsymbol{x}_k &= \boldsymbol{x}_{k-1} + \boldsymbol{\omega}_{k-1} \\ z_k &= g(\boldsymbol{x}_k, k) + \eta_k \end{aligned} \qquad (3)$$

where, as anticipated in the previous Section, the discrete cycle index k is the input of the non-linear MLP mapping $g(\cdot)$ in (2), z_k is the output of $g(\cdot)$, i.e., the Li-ion battery capacity, the random process $\boldsymbol{\omega}_k \in \mathbb{R}^{n_x \times 1}$ is the stochastic component of the MLP model parameter evolution and $\eta_k \in \mathbb{R}$ is the measurement noise [33]. Both random processes are here assumed to be Gaussian. The evolution of the MLP parameters during their updating process is thus modeled as a discrete random walk driven

by a process noise $\boldsymbol{\omega}_k$, large enough to suitably "explore" the parameters' space and to guarantee a certain degree of flexibility so as to possibly adapt to any, possibly unexpected, changes in the process dynamics. Note that the MLP network $g(\boldsymbol{x}_k, k)$ is here taken as the model of the measurement equation, i.e., the relationship between the hidden states $\boldsymbol{x}_k$ and the observations z_k at cycle k.

In this framework, the particle filter is used to recursively estimate the posterior probability density function (pdf) $p(\boldsymbol{x}_k|\boldsymbol{z}_{1:k})$ of the MLP parameters $\boldsymbol{x}_k$, given the set of observations $\boldsymbol{z}_{1:k}$ up to the current k-th cycle. The nonlinearity of the dynamic state-space model in (3), which is due to the representation of the measurement equation by a MLP model, is such that analytical solutions of the optimal prediction-update Bayesian recursion for $p(\boldsymbol{x}_k|\boldsymbol{z}_{1:k})$ cannot be obtained [33], [34]. Thus, similarly to [2], here it is proposed to estimate the posterior pdf $p(\boldsymbol{x}_k|\boldsymbol{z}_{1:k})$ by means of the sampling importance resampling (SIR) PF algorithm. The Monte Carlo estimator of the MLP parameters posterior pdf reads:

$$\hat{p}(\boldsymbol{x}_k|\boldsymbol{z}_{1:k}) \approx \sum_{i=1}^{N_s} w_k^{(i)} \delta\left(\boldsymbol{x}_k - \boldsymbol{x}_k^{(i)}\right) \quad (4)$$

where the importance samples $\boldsymbol{x}_k^{(i)}$, $i = 1, \dots, N_s$ are N_s independent and identically distributed realizations of the system state vector, drawn from the importance pdf $p(\boldsymbol{x}_k|\boldsymbol{x}_{k-1})$ (also called the "prior"). The terms $w_k^{(i)}$, $i = 1, \dots, N_s$ are the normalized importance weights and $\delta(\cdot)$ is the Kronecker delta. The SIR allows for a recursive estimation of the normalized importance weights as [34]:

$$w_k^{(i)} = \frac{\widetilde{w}_k^{(i)}}{\sum_{i=1}^{N_s} \widetilde{w}_k^{(i)}} \quad (5)$$

where the non-normalized weights $\widetilde{w}_k^{(i)}$ are:

$$\widetilde{w}_k^{(i)} = w_{k-1}^{(i)} p\left(z_k|\boldsymbol{x}_k^{(i)}\right); \qquad i = 1, \dots, N_s \quad (6)$$

The function $p\left(z_k|\boldsymbol{x}_k^{(i)}\right)$ is the likelihood of the observation z_k, i.e., the probability of observing z_k given the set of MLP model parameters $\boldsymbol{x}_k^{(i)}$.

A resampling scheme is implemented after the weight normalization, in order to avoid the well-known sample impoverishment problem [33], [34].

In this particle filtering framework, the posterior distribution $p(\boldsymbol{x}_k|\boldsymbol{z}_{1:k})$, where the vector $\boldsymbol{z}_{1:k}$ represents the sequence of capacity observations in correspondence of the cycle sequence $(1:k)$, is estimated at each cycle k, as soon as a new observation of the battery capacity z_k becomes available. The set of samples and associated weights $\left\{\boldsymbol{x}_k^{(i)}, w_k^{(i)}\right\}$ identifies N_s MLP network models, which are here used also to predict the capacity behavior over the future charge-discharge cycles, i.e., the EOL, as it shall be discussed later.

Operatively, the random walk in (3) is assumed to be driven by uncorrelated, zero-mean Gaussian noises, i.e., $\boldsymbol{\omega}_k \sim \mathcal{N}\left(0, \Sigma_{\boldsymbol{\omega}_k}\right)$, where the covariance matrix $\Sigma_{\boldsymbol{\omega}_k}$ is diagonal.

In general, the choice of the noise variances is not an easy task: too small values may hamper a proper (and reasonably fast) exploration of the state-space, whereas too large values do not guarantee a satisfactory state estimation, as discussed by many authors in literature. A common approach for achieving a satisfactory trade-off is that of letting the variances decrease from an initial value as the estimation process progresses, so as to guarantee the convergence of the algorithm[36], [37].

In order to achieve a satisfactory trade-off between the filter speed of convergence and its capability of effectively exploring the state-space, according to the discussion in [2], the following expression for the process noise covariance matrix as a function of the discrete cycle k is here proposed:

$$\Sigma_{\boldsymbol{\omega}_k} = \left(\sigma_0\, e^{-\frac{k}{\sigma_1}} + \sigma_2\right) \cdot \mathrm{I} \qquad (7)$$

where the artificial variances of all the parameters $\boldsymbol{x}_k$ (MLP network weights and biases) are taken to be the same and I is the $n_x \times n_x$ identity matrix. By using this expression for the variances, it is possible to easily and intuitively set their starting value $(\sigma_0 + \sigma_2)$, settling value (σ_2) and rate of decrease $(\frac{1}{\sigma_1})$. The measurement noise η_k is also zero-mean, Gaussian with variance σ_η^2 and independent from k, i.e., $\eta_k \sim \mathcal{N}\left(0, \sigma_\eta^2\right)$.

The likelihood function in (6) is chosen to be:

$$\mathcal{L}_k^{(i)} = p\left(\boldsymbol{z}_{1:k}|\boldsymbol{x}_k^{(i)}\right) =$$
$$= \left((2\pi)^{k+1}\left|\Sigma_\eta\right|\right)^{-0.5} \exp\left\{-\frac{1}{2}\left(\boldsymbol{z}_{1:k} - g\left(\boldsymbol{x}_k^{(i)}, 1:k\right)\right)^T \Sigma_\eta^{-1}\left(\boldsymbol{z}_{1:k} - g\left(\boldsymbol{x}_k^{(i)}, 1:k\right)\right)\right\} \qquad (8)$$

where the terms $g\left(\boldsymbol{x}_k^{(i)}, 1{:}\,k\right)$ are the predictions of the *i-th* MLP network with parameters $\boldsymbol{x}_k^{(i)}$ in correspondence of the input sequence $(1{:}\,k)$. By doing so, the particles weights are computed on the basis of the whole sequence of capacity observations up to the current step *k*, thus increasing the filter robustness [2].

On the basis of the results obtained in [2], in order to improve the efficiency of the approach, a pre-training of the MLP on the basis of a properly chosen degradation curve observed in some reference Li-Ion battery is performed. By doing so, in fact, the network weights are expected to reach a region of the parameter space closer to the optimal one for the observations that will be used during the diagnostic/prognostic task. However, Li-Ion batteries may show very different degradation behaviors, due to the operating conditions, the battery types, etc., so that a pre-training on data very different from those that will actually be used might be even misleading. In order to soften this issue and improve the algorithm flexibility and robustness, the observations in the pre-training set are shifted so that the first capacity observation (corresponding to $k = 1$) is equal to the first available observation of the new degradation dynamics. By so doing, at least the initial values of the observations processed by the filter are going to be similar to those seen in the pre-training stage, so that the algorithm is not forced to adapt to the different values from the very beginning of the estimation process and can, more gradually, adjust itself to the, possibly, different trend. Note also that, as anticipated at the end of Section 2.1, the input cycles and the corresponding capacity observations in the pre-training set are normalized by their means and standard deviations.

At each cycle k the posterior pdf of the end of life time, $p(\mathrm{EOL}_k|\boldsymbol{z}_{1:k})$, is estimated by letting the N_s MLP networks, associated to the N_s parameter samples (particles) $\boldsymbol{x}_k^{(i)},\ i = 1, \dots, N_s$, evolve until their capacity predictions reach the predefined threshold [38]. In order to do so, a vector $[k+1, \dots, k+p]$ (with the number of steps ahead, p, properly chosen so as to be sure that each network prediction reaches the threshold within this time horizon) is fed to each of the N_s MLP networks. Then, the sample posterior pdf $\hat{p}(\tilde{z}_{k+l}|\boldsymbol{z}_{1:k})$ of the capacity predictions at the future cycles l is constructed on the basis of the corresponding MLP outputs, thus obtaining:

$$\hat{p}(\tilde{z}_{k+l}|\boldsymbol{z}_{1:k}) \approx \sum_{i=1}^{N_s} w_k^{(i)} \delta\left(\tilde{z}_{k+l} - g\left(\boldsymbol{x}_k^{(i)}, t_{k+l}\right)\right) \qquad (9)$$

At the same time, the $\mathrm{EOL}_k^{(i)},\ i = 1, \dots, N_s$ obtained in correspondence of each MLP network are used to build the sample posterior pdf of the EOL_k, $\hat{p}(\mathrm{EOL}_k|\boldsymbol{z}_{1:k})$:

$$\hat{p}(\mathrm{EOL}_k|\boldsymbol{z}_{1:k}) \approx \sum_{i=1}^{N_s} w_k^{(i)} \delta\left(\mathrm{EOL}_k - \mathrm{EOL}_k^{(i)}\right) \quad (10)$$

Note that the remaining useful lifetime of the battery at the current cycle k is $\mathrm{RUL}_k = (\mathrm{EOL}_k - k)$. In order to enhance the prognostic capabilities, both in terms of accuracy and stability, improving the strategy suggested in [2], the algorithm introduced above is further modified by empirically enhancing the likelihood function in (8). Operatively, at each cycle k, before the resampling stage, a small number $N_{trivial}$ of particles $\boldsymbol{x}_k^{(i)},\ i = 1, \dots, N_{trivial}$ (note that each particle is a sample of the MLP network parameters) are substituted by a new particle (called hereafter "trivial") built by training the same MLP network architecture with a set made up of the actual sequence of degradation observations $\boldsymbol{z}_{1:k}$ up to the current cycle k and the observations $\boldsymbol{z}_{k+1:end}$ of the initial pre-training set, after properly shifting them according to a procedure similar to that illustrated for the first training of the MLP, so as to smoothly connect the two observation sequences. Figure 2 illustrates how the training datasets for the trivial MLP networks are created at two different cycles for a capacity degradation example (top left and right figures) and the resulting trivial predictions (bottom left and right figures): the blue circles represent the observations of the current Li-ion degradation curve, the green circles represent the portion of the pre-training observations appended to the trivial training dataset, the grey circles are the pre-training dataset and the dashed green lines are examples of particles' trajectories projected in the future at the two cycles considered. As qualitatively verified by the authors, but not shown here for brevity's sake, this scheme is qualitatively shown to improve the prognostic performances with the parameter $N_{trivial}$ set to values ranging between 1 and 10. In fact, a few particles with a good behavior up to (at least) the current cycle k should contribute to maintain the whole particle swarm "close" to the actual degradation dynamics, especially when unexpected behaviors of the observations and/or exceedingly deviated values of the noises occur. On the other hand, using a too large number of identical "trivial" particles would tend to be equivalent to directly using the trivial MLP for the predictions, thus not exploiting the filtering capabilities of the filter.
a sufficient number of trivial particles will speed up the filter convergence and stability, whereas too large values of the parameter $N_{trivial}$ will worsen the filter prediction outside the training domain.

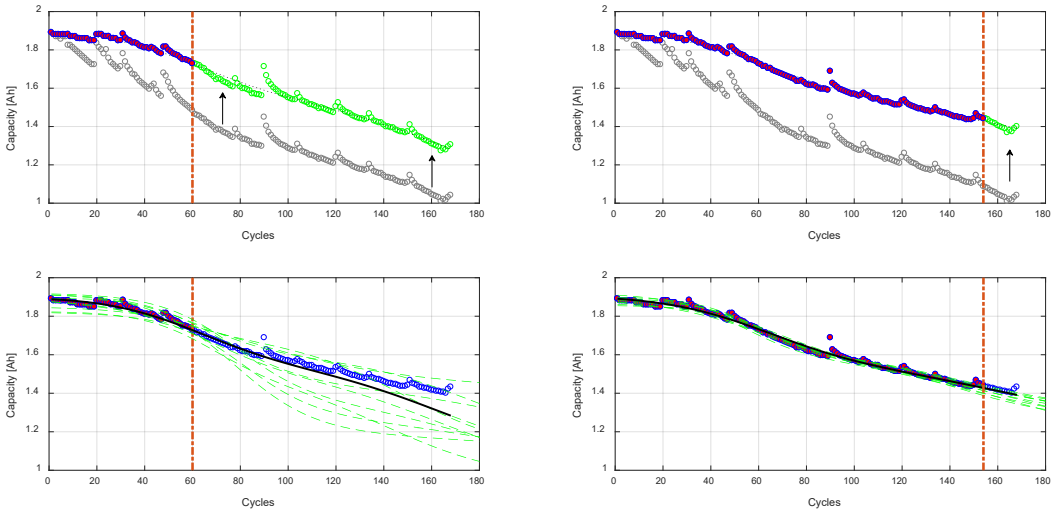

Figure 2: Illustration of the trivial MLP networks creation procedure at two different cycles (marked with red, dashed lines) for an illustrative degradation curve (top figures) and corresponding trivial network predictions (black solid line, bottom figures.

2.3 Log-likelihood ratio (LLR)-based diagnostic module

In this Section, the particle filter-based prognostic algorithm described above is equipped with an additional diagnostic module, based on the estimation of the observations log-likelihood ratio (LLR), for the automatic and on-line detection of any changes in the current dynamic degradation process, due, for example, to variations in the operating conditions (e.g., ambient temperature, charge and discharge procedures, mechanical stresses, etc.) or internal battery failures. To this aim, according to [39], [40] and [41], the following diagnostic index (LLR_k), based on the particle mean likelihoods, is introduced:

$$LLR_k = ln\left(\frac{\frac{1}{N_s}\sum_{i=1}^{N_s}\mathcal{L}_{k-1}^{(i)}}{\frac{1}{N_s}\sum_{i=1}^{N_s}\mathcal{L}_{k}^{(i)}}\right) \quad (11)$$

where the numerator and the denominator are the mean likelihoods of the particles (i.e., the MLP network parameters) at the cycles $k-1$ and k, respectively. Operatively, if the mean likelihood of the particles at cycle k is equal to that at cycle $k-1$, i.e., $\mathcal{L}_k = \mathcal{L}_{k-1}$, then $LLR_k = 0$ and, probably, the observed capacity behavior does not differ significantly from that predicted by the N_s particles, i.e. the N_s MLP networks. On the other hand, if, on average, $\mathcal{L}_k < \mathcal{L}_{k-1}$, then $LLR_k > 0$, meaning

that the observation at k is not as "coherent" as that at $k-1$ with the degradation dynamics predicted by the MLP networks, probably due to some changes in the observed capacity with respect to the expected behavior, occurred between cycles k and $k-1$. The case $\mathcal{L}_k > \mathcal{L}_{k-1}$, so that $LLR_k < 0$, may also occur after a change in the degradation behavior has occurred and when the algorithm eventually "understands", or equivalently adjusts to, the new dynamics: this behavior will be quite clear in the Figures described in the next Section. As a matter of fact, even simple statistical fluctuations in the capacity observations may give rise to situations where, on average, $\mathcal{L}_k < \mathcal{L}_{k-1}$: thus, in order to build a robust diagnostic module, an anomaly is assumed to be detected any time the index in (11) is larger than some threshold. As well known in literature, setting such a threshold is not, in general, an easy task, being necessary to identify a proper trade-off between the probability of false alarms (threshold too low) and the probability of missed alarms (threshold too high). In order to avoid the problem of the identification of the optimal trade-off, a dynamic threshold as $LLR_{th,k} = 3 \cdot \bar{\sigma}(LLR_{1:k})$ is set, where $\bar{\sigma}(LLR_{1:k})$ is the standard deviation of the log-likelihood ratios up to the current cycle k. By so doing, the threshold is automatically and dynamically set to a value that is typically used for triggering alarms in diagnostic applications [42]. Note that, in correspondence of events like recharge or storage, the Li-ion batteries may partially recover their capacity, thus resulting in short increasing trends in the, otherwise, generally decreasing capacity measurements. These behaviors could mislead the log-likelihood ratio-based diagnostic tool, since they are interpreted as anomalies of the expected degradation dynamics. In order to overcome this problem, a simple modification of the algorithm is implemented, which relies on the evaluation of the average derivative of the observations over the last few cycles (say 3 cycles). When the threshold $LLR_{th,k}$ is overcome, if the average observation derivative is positive (meaning that the capacity has actually improved), then i) the alarm is not triggered, ii) $\bar{\sigma}(LLR_{1:k})$ is not updated and iii) the LLR_k is re-initialized to 0, until the observations become smaller than the last, pre-anomaly one. In fact, keeping $\bar{\sigma}(LLR_{1:k})$ equal to the same value reached before the (positive) anomaly occurred (instead of letting it increase) allows to avoid that future anomalies are not detected (missed alarms) due to too large threshold values. At the same time, by forcing the LLR_k to zero, a sequence of false alarms, due to the fact that a few cycles are needed by the battery capacity to go back to the same pre-anomaly levels, can be avoided. These behaviors are further explained in Figure 3 with reference to an illustrative capacity degradation example: the top Figure shows the reference capacity degradation curve, the middle Figure shows the $\boldsymbol{LLR_k}$ (blue, solid line) and the corresponding $\pm\bar{\boldsymbol{\sigma}}(\boldsymbol{LLR_{1:k}})$ (red, dashed lines), the bottom Figure displays the modified $\boldsymbol{LLR_k}$ (blue, solid line) and the corresponding lower $\pm\bar{\boldsymbol{\sigma}}(\boldsymbol{LLR_{1:k}})$ (red, dashed lines) compared to the previous one (grey dashed lines).

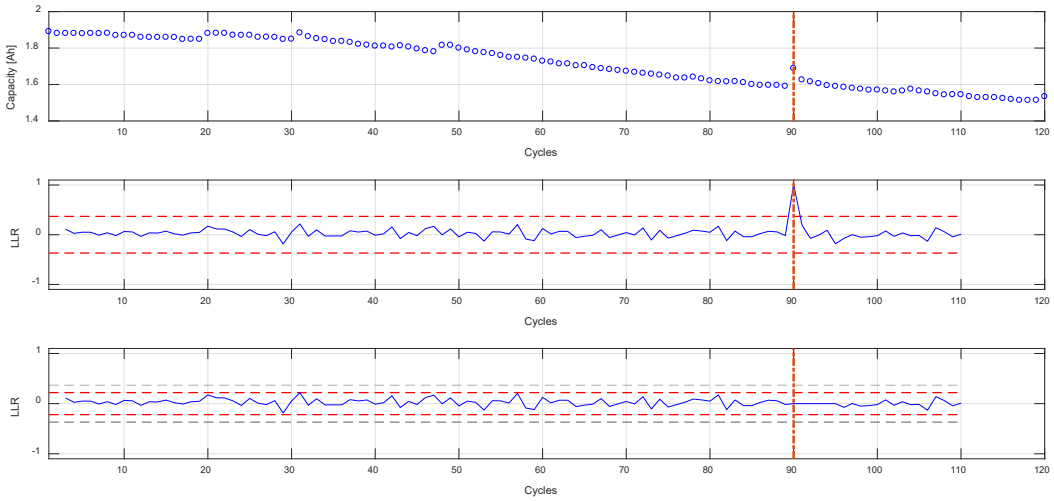

Figure 3: Procedure adopted to avoid triggering detection alarms in correspondence of capacity gains, shown with reference to the same illustrative example of Figure 2.

3. RESULTS

In this Section the prognostic and diagnostic capabilities of the proposed algorithm are demonstrated, with reference to the SOL datasets by the Prognostic Center of Excellence at NASA Ames Research Center and of the CALCE Battery Group ([29] and [30]), shown in Figure 4. In particular, the datasets labeled "NASA 1, 2 and 3" are taken from the NASA Ames research center database, whereas the one labeled "CALCE" is taken from the CALCE database. Note that these datasets are obtained under controlled laboratory conditions and at constant discharge rates (constant current).

According to the initialization procedure illustrated in Section 2, for convenience, but with no loss of generality, NASA 2 is chosen as the pre-training dataset for the MLP networks, as it shows the broader range of variation of its capacity. The PF-MLP algorithm is run with $N_S = 500$ particles; a larger number of particles would excessively slow down the algorithm, with no benefits in terms of mean of the RUL posterior pdf $\hat{p}(\mathrm{EOL}_k|\boldsymbol{z}_{1:k})$, as qualitatively verified by the authors. The parameters defining the process noise variance in (7), chosen on the basis of a trial and error procedure, are: $\sigma_0 = 5 \cdot 10^{-3}$, $\sigma_1 = 10^2$ and $\sigma_2 = 10^{-4}$. The standard deviation of the measurement noise is taken equal to $\sigma_\eta = 10^{-1}$. Note only that the factor σ_1 is taken large enough for the process variance to decrease quite slowly with respect to the degradation times, so as to avoid a too large reduction of the parameter space spanned by the particles at later times.

In practical applications, Li-ion batteries are usually assumed to be failed when their capacity drops below 80% of its initial value. However, in the case studies shown in this work, the failure threshold will be set in order to both maximize the number of available capacity observations to be processed by the algorithm and enhance the readability of the results, with no loss of generality.

3.1 MLP-PF for the EOL prognosis of Li-Ion batteries

First, the algorithm is tested in an ideal situation, i.e., for predicting the RUL of NASA 2 after the initialization of the MLPs of all the particles is performed using the same NASA 2 dataset. Indeed, in this case it is expected that the best algorithm performance, since the MLP parameters already start from good, optimized values. The failure threshold is here set to a value that is 2% larger than the lowest capacity observation available in the dataset. Figure 5 shows that, after an initial, rather short, adaptation period, the actual RUL (grey, dashed line) is always between the 5^{th} and 95^{th} percentiles of the estimated RUL posterior distribution (green lines) and rather close to the RUL posterior mean (red line). The initial deviation from the actual RUL, which should not be expected since all the MLP particles are trained on the same NASA 2 degradation trajectory, is actually due to the fact that, the MLPs input cycles are here normalized with respect to a mean sufficiently larger than that of the pre-training set: in fact, since the RUL of a new battery is not a priori known, the prediction performances of the MLPs are better if the inputs are normalized over an interval which includes the unknown, current one. Hence, the first RUL predictions tend to be larger than the actual RUL (~160 cycles). This approach will be adopted also for all the following examples.

Then, the algorithm is applied to predict the RUL of NASA 1 (1850mAh initial capacity). Figure 6 shows that the prognostic performances are satisfactory, although the actual RUL (grey, dashed line) is systematically slightly underestimated by the mean of the RUL poterior distribution (red lines). This is motivated by the fact that the degradation curve of NASA 1 (as also those of the other types of battery considered in this work, see Figure 4) shows several positive spikes due to the performances recovery phenomenon [43]. This behavior is actually beneficial to the SOL of the battery, thus leading, in general, to larger RULs. However, these "anomalies" cannot be predicted by the algorithm, which, nevertheless, is shown to be capable of adapting to the changed degradation dynamics and, consequently, to quickly update the RUL posterior pdf estimate.

The batteries used in the previous tests show rather similar trends, even if belonging to different types (NASA 1 1850mAh, NASA 2 2000mAh). So, according to what illustrated in Section 2 with regards to the pre-training procedure, satisfactory results are to be expected. However, as anticipated in the Introduction, the aim of this work is that of developing a flexible computational prognostic/diagnostic

tool capable of automatically dealing with different types of Li-Ion batteries. In order to demonstrate this capability, now the prognostic performances of the algorithm when predicting the RUL of a very different type of battery, i.e. the 1100mAh battery (CALCE) taken from the CALCE database [30], are presented. The comparison of the training dataset (NASA 2 from the NASA Ames database) and that used for testing the algorithm (CALCE from the CALCE database) in Figure 4 shows how the two datasets differ both in terms of absolute capacities and lifetimes. In this case, the failure threshold is set to a value that is 20% less of the initial maximum battery capacity. Thus, the pre-training dataset normalization procedure illustrated in Section 2.2 plays an important role for guaranteeing a fast adaptation of the algorithm to the new degradation dynamics. In spite of the fluctuating appearance of the RUL estimate, Figure 7 shows that the RUL predictions (red lines) are still quite satisfactory and that the MLP-PF behaves as expected. In fact, at earlier times, the capacity observations are still rather close to those of the normalized ones used for the initial MLPs training, so that all the particles (i.e., the MLP networks parameters) do not significantly differ from the training ones, and the corresponding predictions are coherent with the normalized training RUL trajectory (blue, dashed line). As soon as the algorithm starts perceiving the different behavior of the actual observations, at approximately 100 cycles, it suddenly changes the RUL prediction, which accordingly becomes much higher (the true RUL is shown by the grey, dashed line). The weird behavior of the 5^{th} percentile (green line) is due to the presence of the "trivial" particles (see Section 2.2), which tend to adapt more slowly to the new degradation dynamics, thus also being responsible for the initial overestimation of the RUL after approximately cycle 100.

Note that, as verified by the authors, but not reported here for brevity's sake, the filter is capable of converging even if no pre-training sets are available, provided the degradation process is not too fast, since the adaptation phase generally requires longer times.

3.2 MLP-PF diagnostic module

Now the diagnostic capabilities of the proposed MLP-PF algorithm, equipped with the diagnostic module described in Section 2.3, are tested. The goal is that of detecting anomalies occurring during the "normal" degradation process of the batteries while, simultaneously, performing prognosis.

To this aim, first a possible sudden drop of the battery capacity is considered, which is due, for example, to some problems during the battery recharge stage or to changed storage conditions. Since the databases available do not present such anomalies, being obtained in controlled laboratory conditions at constant discharge currents, the drop is artificially introduced by manipulating the available data. Figure 8 (a) shows the manipulated observations with the artificial drop (5% of the

current capacity) introduced at cycle 120 of the NASA 3 dataset. As, for the previous case studies, the MLP particles are pre-trained using the properly normalized NASA 2 dataset, according to the procedure described in Section 2.2. The failure threshold is set to a value that is 2% larger than the lowest capacity observation available in the dataset. Figure 8 (b) shows the corresponding RUL prediction. The anomaly leads to a reduction of the actual RUL, which decreases from 156 to 125 cycles, which is not fully captured by the algorithm at earlier times: the RUL estimate should, in fact, slightly overestimate the actual one, since the anomaly still has to occur. However, the magnitude of the drop is rather small, especially if compared to that of several performances recovery phenomena experienced by the battery during its life, and its late occurrence does not significantly modify the battery RUL. Nevertheless, the LLR of (11) and the corresponding anomaly indicator, shown in Figure 8 (c) and (d), respectively, flawlessly and timely capture the change in the degradation dynamics. Note also that, according to the procedure illustrated in Section 2.3, all the beneficial performances recovery phenomena are not detected by the anomaly indicator. On the other hand, the capacity recovery occurring at approximately 90 cycles is so large in magnitude that the algorithm correctly starts changing its RUL prediction, which becomes larger than the actual one.

Then, the MLP-PF performances are tested with respect to a different, more subtle, type of anomaly. More specifically, with reference again to the NASA 3 dataset, an increase of the degradation speed is considered, as illustrated in Figure 9 (a), which is obtained by artificially modifying the original capacity observations so that their average slope is significantly increased. As expected, the prognostic behavior of the MLP-PF algorithm is very similar to that of the previous case study. Even though the anomaly is now less severe than in the previous case study, the algorithm is still effectively capable of detecting its occurrence, as shown in Figure 9 (c) and (d). The only difference lies in the fact that a short sequence of multiple alarms is triggered after the first detection, which is probably due to the fact that, since the anomaly is less severe, the algorithm takes some additional time to adapt to the new degradation behavior. These multiple alarms are observed also in other situations, not reported here for brevity's sake: however, since they always occur in compact sequences, we believe they do not pose significant problems to the detection task, leaving the solution to future work. Moreover, a small detection delay can also be observed, which is due to the fact that the anomaly is less severe than the drop of the previous case study, so that a larger number of cycles is required by the filter to reach convergence on the new degradation dynamics.

In both the previous case studies, it is worth underlying how the algorithm tries to adapt to the new degradation dynamics after the anomaly occurs, although not enough cycles are available for the algorithm to converge on the new RUL, due to the anomaly being very close to the EOL. In order to further investigate the algorithm adaptation capabilities, a new case study, based on the CALCE

battery dataset, is considered where an anomaly in the degradation speed is introduced at some earlier battery cycle. In this case, the failure threshold is set to a value that is 35% less of the initial maximum battery capacity.

The prognostic behavior of the algorithm is similar to that of Figure 7. At earlier cycles the RUL estimate is strongly affected by the initial pre-training of the particles' MLPs. Then, at approximately cycle 85, the algorithm suddenly "understands" the actual degradation dynamics and, consequently, predicts a RUL of approximately 250 cycles (i.e., EOL ≅ 350 cycles). After a few cycles (at 100 cycles) the anomaly occurs: the algorithm effectively captures the change of degradation dynamics at approximately 120 cycles (Figure 10 (c) and (d)), the delay being already explained for the previous case study. Accordingly, at the same number of cycles, the predictions are quickly updated in order to match the new dynamic behavior of the battery capacity, thus clearly showing the adaptation capabilities of the proposed MLP-PF filter.

Finally, it is worth noting that the behaviors that are here assumed as "anomalies", can be considered, as done in our tests, as the effects of some failures, but can be also the result of the changing boundary conditions experienced Li-ion batteries during their normal operation in real applications, i.e., as a representative example, when dealing with different discharge rates due to the different current levels required. For this reason, although the original data used in this work refer to capacity fades at constant currents, the tests of the algorithm with the artificial modifications show that the MLP-PF can possibly be applied to more general, realistic cases of varying operating conditions, by exploiting its significant adaptation capability.

xxxxxxx

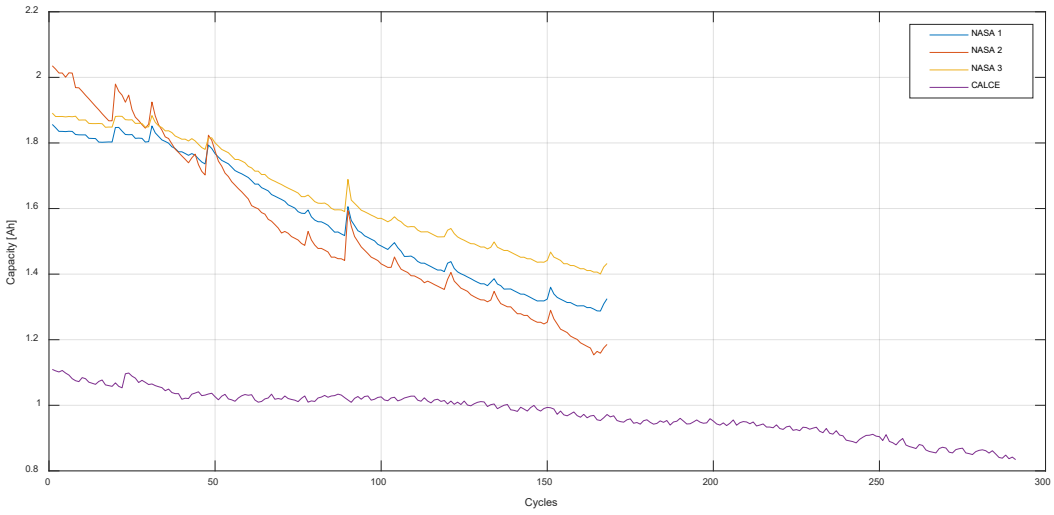

Figure 4: Prognostic Center of Excellence at NASA Ames Research Center [29] and CALCE [30] Battery Group SOL datasets.

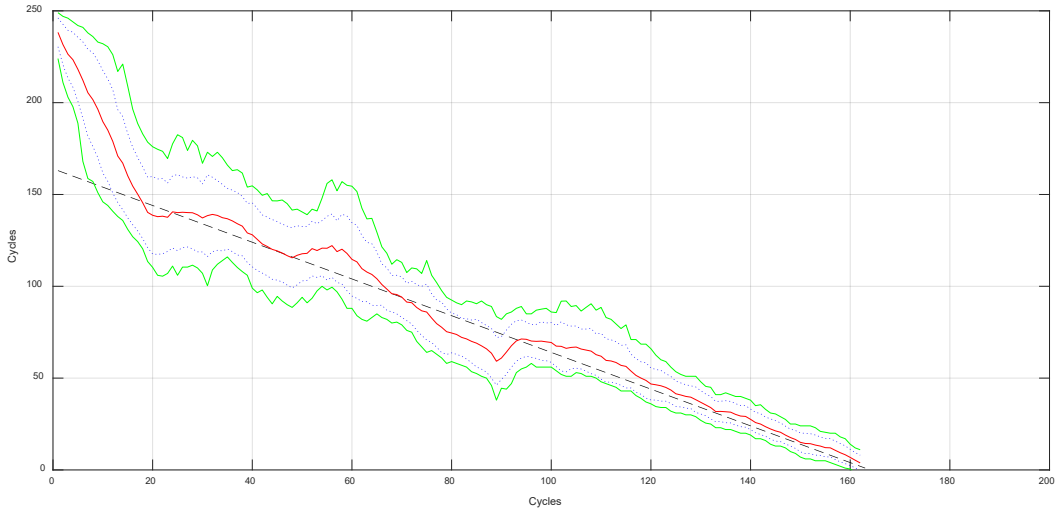

Figure 5: Estimate of the RUL of NASA 2: truth (grey, dashed), mean (red), ± 1 standard deviation (grey dotted) and 5th and 95th percentiles (green) of the RUL posterior pdf.

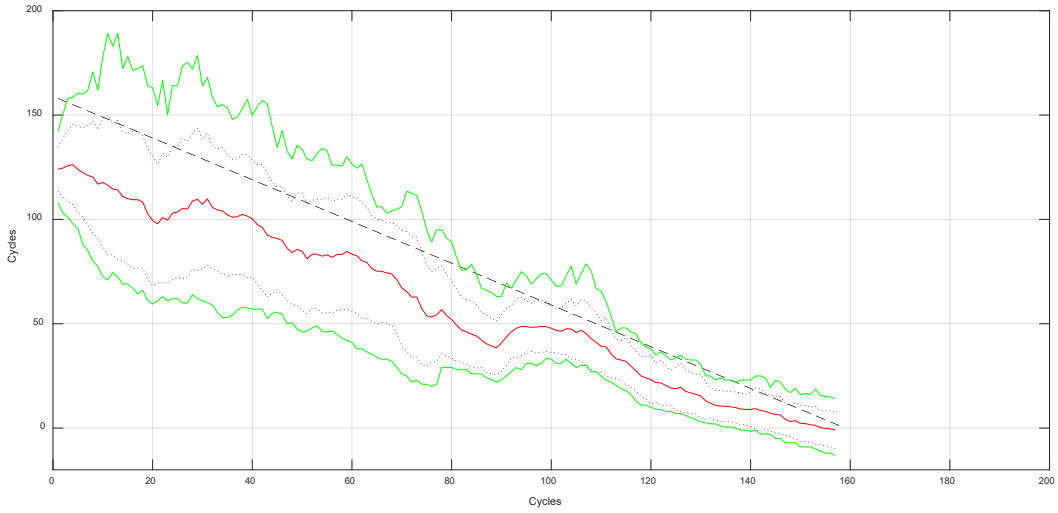

Figure 6: Estimate of the RUL of NASA 1: truth (grey, dashed), mean (red), ± 1 standard deviation (grey dotted) and 5th and 95th percentiles (green) of the RUL posterior pdf.

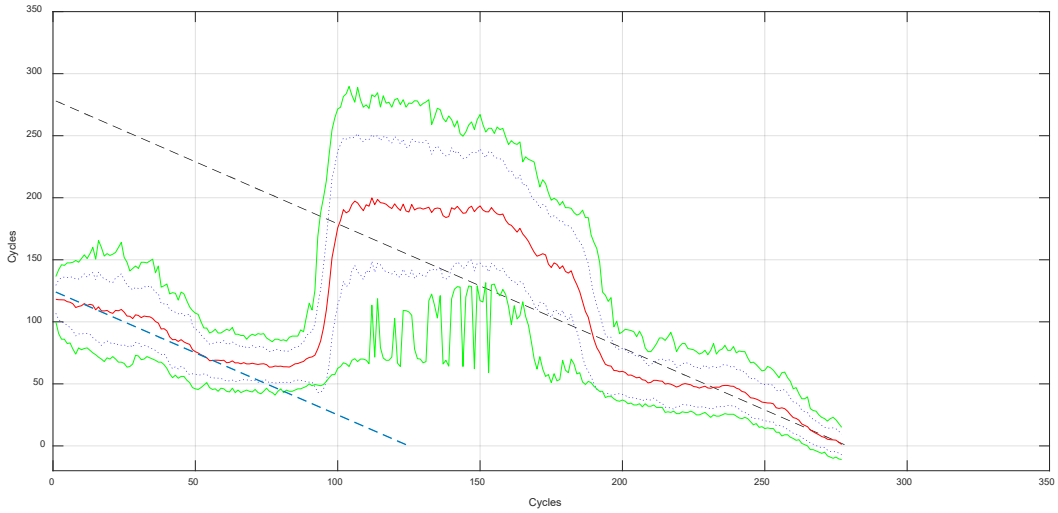

Figure 7: Estimate of the RUL of CALCE: truth (grey, dashed), normalized NASA 2 true RUL (blue, dashed), mean (red), ±1 standard deviation (grey dotted) and 5th and 95th percentiles (green) of the RUL posterior pdf.

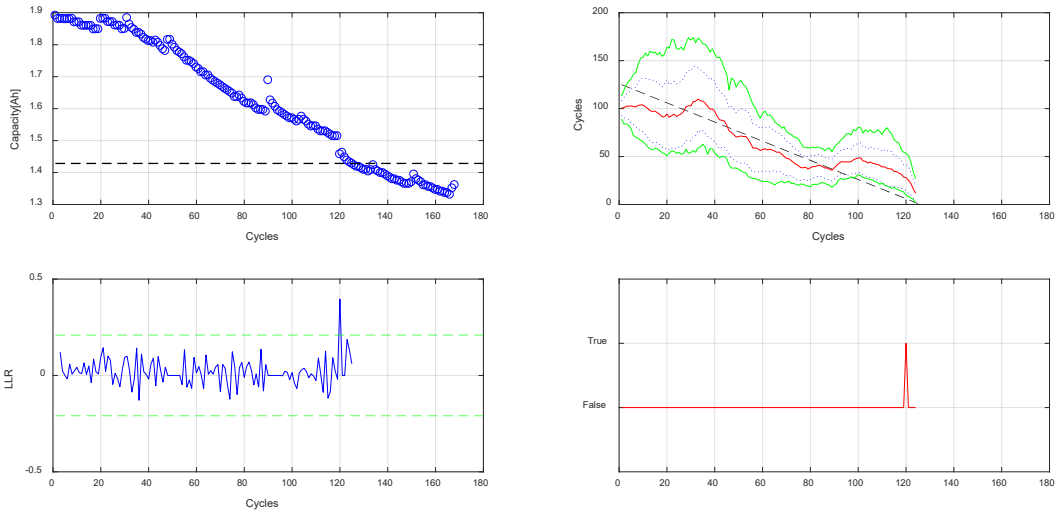

Figure 8: (a) NASA 3 dataset with the artificial drop (5%) at cycle 120; (b) corresponding RUL estimate by the MLP-PF algorithm; (c) log-(mean) likelihood ratio LLR_k (11); (d) anomaly index.

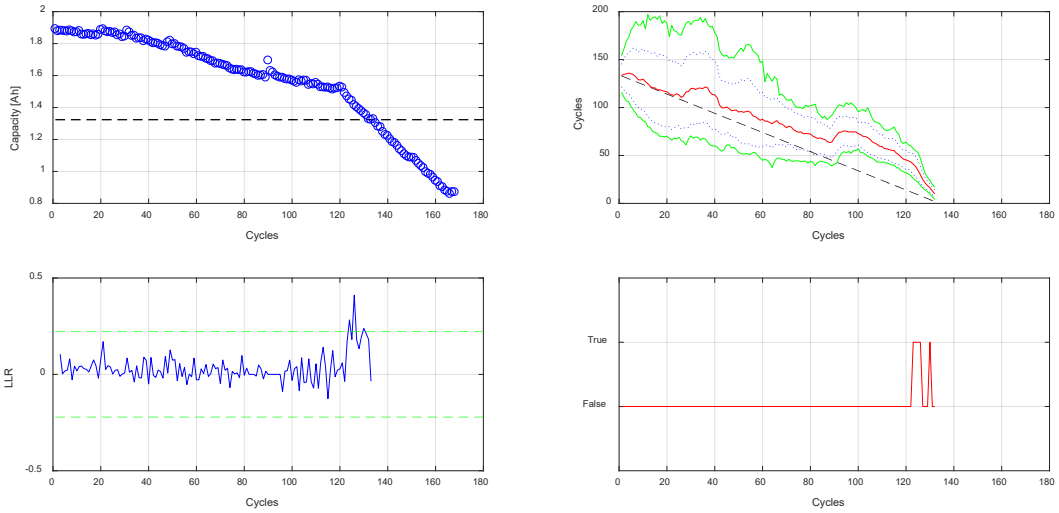

Figure 9: (a) NASA 3 dataset with the artificial slope change at cycle 120; (b) corresponding RUL estimate by the MLP-PF algorithm; (c) log-(mean) likelihood ratio LLR_k (11); (d) anomaly index.

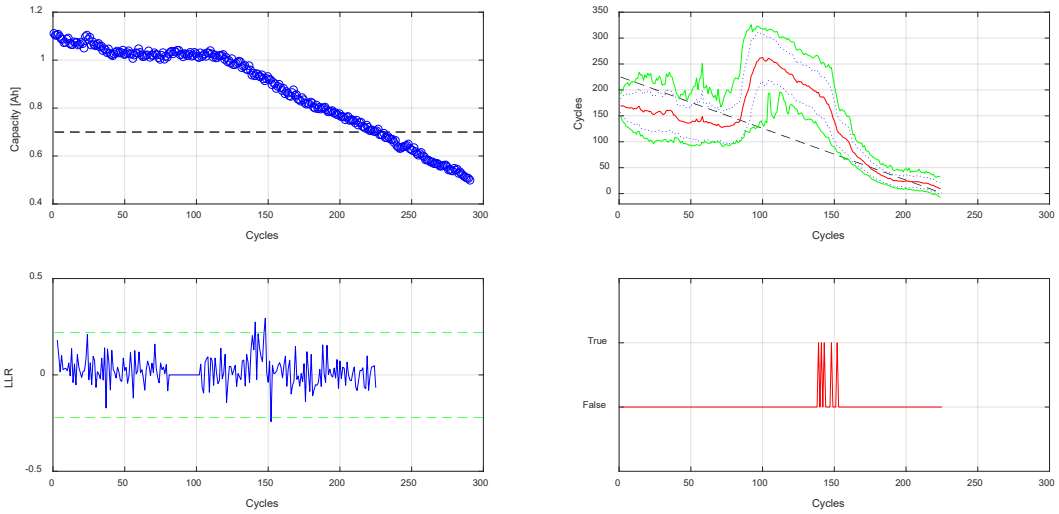

Figure 10: (a) CALCE dataset with the artificial slope change at cycle 120; (b) corresponding RUL estimate by the MLP-PF algorithm; (c) log-(mean) likelihood ratio LLR_k (11); (d) anomaly index.

4. CONCLUSIONS

In this work an original adaptation of an algorithm introduced by some of the same authors for performing adaptive and on-line prognosis of the EOL of Li-Ion batteries and diagnosis of their SOH have been proposed.

The use of a MLP neural network for approximating the observation equation significantly contributes to the adaptability of the prognostic/diagnostic tool adopted. As demonstrated by the results of this work, the advantage of this feature is twofold. First, the algorithm can be effectively applied to many different types of batteries, with very different degradation behaviors, with no need to derive *ad hoc* physics-based models for any possible application. Second, the algorithm is capable of automatically re-configuring itself in case a change of the degradation dynamics occurs, thus still guaranteeing the effectiveness of the prognostic and diagnostic tasks. At the same time, the on-line training strategy offered by the particle filter framework provides the neural network with the capability of automatically adapting to different degradation behaviors in real time, thus avoiding the need for gathering huge amounts of off-line experimental data to cover all possible operative/fault conditions.

Finally, although no accurate physics-based models are used, the proposed LLR-based diagnostic index, along with a definition of a suitable adaptive detection threshold, have shown to be effective in detecting the cycle at which the degradation behavior undergoes changes in the dynamics, possibly due to faults and anomalies.

A possible disadvantage related to the use of adaptive surrogate models lies, on the other hand, in the response times offered by the proposed algorithm (i.e., the number of cycles required to react to changes in the degradation dynamics), which, on average, tend to be larger than those required by similar methods relying on the use of physics-based models. Indeed, a choice must be made between the two types of approaches, mainly depending on the objectives of the application under analysis and the quantity of information available (observations, models, etc.).

An important future activity is a systematic analysis of the relative importance of the different algorithm input parameters (often qualitatively identified in this work on the basis of trial and error procedures) in affecting the algorithm performances. Indeed, this would require to resort also to properly defined prognostic performance measures, in order to be able to quantify the effects of the different input parameters. On the basis of the available scientific literature and our experience, the process and measurement noise parameters would, most likely, be confirmed as the most critical for achieving satisfactory performances in both the diagnostic and prognostic tasks.

5. REFERENCES

[1] K. Goebel, B. Saha, and P. Christophersen, “Prognostics in Battery Health Management,” *IEEE Instrum. Meas. Mag.*, pp. 33–40, 2008.

[2] C. Sbarufatti, M. Corbetta, M. Giglio, and F. Cadini, “Adaptive prognosis of lithium-ion batteries based on the combination of particle filters and radial basis function neural networks,” *J. Power Sources*, vol. 344, pp. 128–140, 2017.

[3] M. B. Pinson and M. Z. Bazant, “Theory of SEI Formation in Rechargeable Batteries: Capacity Fade, Accelerated Aging and Lifetime Prediction,” *J. Electrochem. Soc.*, vol. 160, no. 2, pp. A243–A250, 2012.

[4] W. He, N. Williard, M. Osterman, and M. Pecht, “Prognostics of lithium-ion batteries based on Dempster-Shafer theory and the Bayesian Monte Carlo method,” *J. Power Sources*, vol. 196, no. 23, pp. 10314–10321, 2011.

[5] M. Berecibar, I. Gandiaga, I. Villarreal, N. Omar, J. Van Mierlo, and P. Van Den Bossche, “Critical review of state of health estimation methods of Li-ion batteries for real applications,” *Renew. Sustain. Energy Rev.*, vol. 56, pp. 572–587, 2016.

[6] L. Wu, X. Fu, and Y. Guan, “Review of the Remaining Useful Life Prognostics of Vehicle Lithium-Ion Batteries Using Data-Driven Methodologies,” *Appl. Sci.*, vol. 6, no. 6, p. 166, 2016.

[7] J. Guo, Z. Li, and M. Pecht, “A Bayesian approach for Li-Ion battery capacity fade modeling and cycles to failure prognostics,” *J. Power Sources*, vol. 281, pp. 173–184, 2015.

[8] S. Shriram and E. W. Ralph, “State of charge estimation using an unscented filter for high power lithium ion cells,” *Int. J. energy Res.*, vol. 34, pp. 152–163, 2010.

[9] M. Doyle, T. F. Fuller, and J. Newmann, “Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell,” *J. Electrochem. Soc.*, vol. 140, no. 6, pp. 1526–1533, 1993.

[10] M. Safari, M. Morcrette, A. Teyssot, and C. Delacourt, “Multimodal Physics-Based Aging Model for Life Prediction of Li-Ion Batteries,” *J. Electrochem. Soc.*, vol. 156, no. 3, pp. A145–A153, 2009.

[11] G. Ning, B. Haran, and B. N. Popov, “Capacity fade study of lithium-ion batteries cycled at high discharge rates,” *J. Power Sources*, vol. 117, no. 1–2, pp. 160–169, 2003.

[12] Y. Zou, X. Hu, H. Ma, and S. E. Li, “Combined State of Charge and State of Health estimation over lithium-ion battery cell cycle lifespan for electric vehicles,” *J. Power Sources*, vol. 273, pp. 793–803, 2015.

[13] C. Hu, G. Jain, P. Tamirisa, and T. Gorka, “Method for estimating capacity and predicting

remaining useful life of lithium-ion battery," *Appl. Energy*, vol. 126, pp. 182–189, 2014.

[14] M. Ye, H. Guo, and B. Cao, "A model-based adaptive state of charge estimator for a lithium-ion battery using an improved adaptive particle filter," *Appl. Energy*, vol. 190, pp. 740–748, 2017.

[15] S. B. Peterson, J. Apt, and J. F. Whitacre, "Lithium-ion battery cell degradation resulting from realistic vehicle and vehicle-to-grid utilization," *J. Power Sources*, vol. 195, no. 8, pp. 2385–2392, 2010.

[16] G. Jin, D. E. Matthews, and Z. Zhou, "A Bayesian framework for on-line degradation assessment and residual life prediction of secondary batteries in spacecraft," *Reliab. Eng. Syst. Saf.*, vol. 113, no. 1, pp. 7–20, 2013.

[17] D. Liu, J. Pang, J. Zhou, Y. Peng, and M. Pecht, "Prognostics for state of health estimation of lithium-ion batteries based on combination Gaussian process functional regression," *Microelectron. Reliab.*, vol. 53, no. 6, pp. 832–839, 2013.

[18] D. Wang, Q. Miao, and M. Pecht, "Prognostics of lithium-ion batteries based on relevance vectors and a conditional three-parameter capacity degradation model," *J. Power Sources*, vol. 239, pp. 253–264, 2013.

[19] J. Wu, C. Zhang, and Z. Chen, "An online method for lithium-ion battery remaining useful life estimation using importance sampling and neural networks," *Appl. Energy*, vol. 173, pp. 134–140, 2016.

[20] M. Broussely, S. Herreyre, P. Biensan, P. Kasztejna, K. Nechev, and R. J. Staniewicz, "Aging mechanism in Li ion cells and calendar life predictions," *J. Power Sources*, vol. 97–98, pp. 13–21, 2001.

[21] E. V. Thomas, I. Bloom, J. P. Christophersen, and V. S. Battaglia, "Statistical methodology for predicting the life of lithium-ion cells via accelerated degradation testing," *J. Power Sources*, vol. 184, no. 1, pp. 312–317, 2008.

[22] R. G. Jungst *et al.*, "Accelerated calendar and pulse life analysis of lithium-ion cells," *J. Power Sources*, vol. 119–121, pp. 870–873, 2003.

[23] Y. Chang, H. Fang, and Y. Zhang, "A new hybrid method for the prediction of the remaining useful life of a lithium-ion battery," *Appl. Energy*, vol. 206, pp. 1564–1578, 2017.

[24] G. Bai, P. Wang, C. Hu, and M. Pecht, "A generic model-free approach for lithium-ion battery health management," *Appl. Energy*, vol. 135, pp. 247–260, 2014.

[25] M. Charkhgard and M. Farrokhi, "State-of-charge estimation for lithium-ion batteries using neural networks and EKF," *IEEE Trans. Ind. Electron.*, vol. 57, no. 12, pp. 4178–4187, 2010.

[26] N. Daroogheh, A. Baniamerian, N. Meskin, and K. Khorasani, "A hybrid prognosis and health

monitoring strategy by integrating particle filters and neural networks for gas turbine engines," in *IEEE Conference on Prognostics and Health Management: Enhancing Safety, Efficiency, Availability, and Effectiveness of Systems Through PHAf Technology and Application, PHM 2015*, 2015.

[27] J. Vetter *et al.*, "Ageing mechanisms in lithium-ion batteries," *J. Power Sources*, vol. 147, no. 1–2, pp. 269–281, 2005.

[28] D. Wang, F. Yang, Y. Zhao, and K.-L. Tsui, "Battery remaining useful life prediction at different discharge rates," *Microelectron. Reliab.*, vol. 78, pp. 212–219, 2017.

[29] B. Saha and K. Goebel, "Battery Data Set. NASA Ames Prognostics Data Repository," *NASA Ames Research Center, Moffett Field, CA*, 2007. [Online]. Available: https://ti.arc.nasa.gov/tech/dash/pcoe/prognostic-data-repository/#battery.

[30] M. Pecht, "Battery Data Set. CALCE," *CALCE Battery Research Group, Maryland, MD*, 2017. [Online]. Available: https://web.calce.umd.edu/batteries/index.html.

[31] J. F. de Freitas, M. Niranjan, A. H. Gee, and A. Doucet, "Sequential Monte Carlo methods to train neural network models," *Neural Comput.*, vol. 12, pp. 955–993, 2000.

[32] C. M. Bishop, *Neural networks for pattern recognition*. 1995.

[33] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Stat. Comput.*, vol. 10, pp. 197–208, 2000.

[34] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/nongaussian bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, 2002.

[35] I. T. Nabney, *NETLAB: Algorithms for Pattern Recognition*. 2002.

[36] S. Schwunk, N. Armbruster, S. Straub, J. Kehl, and M. Vetter, "Particle filter for state of charge and state of health estimation for lithium-iron phosphate batteries," *J. Power Sources*, vol. 239, pp. 705–710, 2013.

[37] N. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," in *IEEE Radar and Signal Processing*, 1993, vol. 140, no. 2, pp. 107–113.

[38] F. Cadini, E. Zio, and D. Avram, "Monte Carlo-based filtering for fatigue crack growth estimation," *Probabilistic Eng. Mech.*, vol. 24, pp. 367–373, 2009.

[39] P. Li and V. Kadirkamanathan, "Particle Filtering Based Likelihood Ratio Approach to Fault Diagnosis in Nonlinear Stochastic Systems," *IEEE Trans. Syst. Man. Cybern.*, vol. 31, no. 3, pp. 337–343, 2001.

[40] F. Cadini, E. Zio, and G. Peloni, "Particle filtering for the detection of fault onset time in hybrid

dynamic systems with autonomous transitions," *IEEE Trans. Reliab.*, vol. 61, no. 1, pp. 130–139, 2012.

[41] J. Grana, D. Wolpert, J. Neil, D. Xie, T. Bhattacharya, and R. Bent, "A likelihood ratio anomaly detector for identifying within-perimeter computer network attacks," *J. Netw. Comput. Appl.*, vol. 66, pp. 166–179, 2016.

[42] Rodríguez-López, López-González, López-Ochoa, and Las-Heras-Casas, "Methodology for detecting malfunctions and evaluating the maintenance effectiveness in generator bearings in wind turbines using generic versus specific models from SCADA data," *Energies*, 2018.

[43] A. Eddahech, O. Briat, and J.-M. Vinassa, "Lithium-ion battery performance improvement based on capacity recovery exploitation," *Electrochim. Acta*, vol. 114, pp. 750–757, Dec. 2013.

[a] "Fuel economy of hybrid electric flight" (ELECTRICAL AIRCRAFT)

[b] "Lithium and lithium ion batteries for applications in microelectronic devices: A review" (MICROELECTRONICS DEVICES)

[c] "Technical and economic assessment of the secondary use of repurposed electric vehicle batteries in the residential sector to support solar energy" (ENERGY GRID)

[d] "The lithium-ion battery: State of the art and future perspectives" (ENERGY GRID)

[e] "Overview of grid connected renewable energy based battery projects in USA" (ENERGY GRID)

[f] "Analyzing system safety in lithium-ion grid energy storage" (ENERGY GRID)

[g] "Can Li-Ion batteries be the panacea for automotive applications?" (AUTOMOTIVE)

[h] "Li ion batteries for aerospace applications" (AEROSPACE)

[i] "Modeling and Simulation of Lithium-Ion Batteries from a Systems Engineering Perspective"