Journal of Power Sources 196 (2011) 4061–4066

Contents lists available at ScienceDirect

Journal of Power Sources

journal homepage: www.elsevier.com/locate/jpowsour

State-of-charge prediction of batteries and battery–supercapacitor hybrids using artificial neural networks

T. Weigert[a,*], Q. Tian[b], K. Lian[b]

[a] Department of Computer Science, Missouri University of Science and Technology, 1870 Miner Circle, Rolla, MO 65401, USA
[b] Department of Materials Science and Engineering, University of Toronto, 184 College Street, Toronto, ON M5S 3E4, Canada

ARTICLE INFO

Article history:
Received 16 September 2010
Received in revised form 24 October 2010
Accepted 26 October 2010
Available online 5 November 2010

Keywords:
State-of-charge prediction
Battery lifetime prediction
Battery–supercapacitor hybrid
Pulse discharge
Artificial neural networks

ABSTRACT

The state-of-charge (SOC) of batteries and battery–supercapacitor hybrid systems is predicted using artificial neural networks (ANNs). Our technique is able to predict the SOC of energy storage devices based on a short initial segment (less than 4% of the average lifetime) of the discharge curve. The prediction shows good performance with a correlation coefficient above 0.95. We are able to improve the prediction further by considering readily available measurements of the device and usage. The prediction is further shown to be resilient to changes in operating conditions or physical structure of the devices.

1. Introduction

There is a pressing need in energy storage to meet the ever increasing demands for sustainable energy. Batteries, supercapacitors, and their hybrid combination are all significant contributing technologies. Hybrid systems containing a battery and a supercapacitor have been experimentally demonstrated to exhibit longer operating times when compared to systems with batteries alone under repetitive high load and high current pulse conditions [1–7].

Electrochemical capacitors (supercapacitors) offer high power density when compared to battery systems and also have a relatively large energy density compared to conventional capacitors [8–11]. Batteries, such as lithium ion batteries, have a high energy density of about $10^5\,\mathrm{J\,kg^{-1}}$. Their power density, however, is low – around $100\,\mathrm{W\,kg^{-1}}$. As a result, conventional batteries cannot meet high power demands when discharged at high currents [10,11]. Combining a supercapacitor in parallel with a battery into a hybrid system broadens the applicability of batteries to higher discharge rates due to the high power density of supercapacitors [1,6,7]. Such hybrid energy storage devices are more efficient than a battery in supplying the total power for use in digital cellular phones, space communications, power distribution systems, uninterrupted power supplies, electric and hybrid vehicles, portable computers, and military applications [7]. In many of these applications, the loads are not constant but rather span a range of power levels. Coupling supercapacitors with batteries is more beneficial under pulsed power loads, which are frequently encountered in communication systems [1].

Being able to predict the lifetime of a power storage device is of immense technical and commercial importance when planning power systems, selecting the most suitable batteries, determining the operating conditions, and planning replacement intervals for batteries [12]. Important questions are how long will a battery of a specific type and manufacturer last under certain operating conditions, and how will the lifetime be affected if the operating conditions are changed?

The lifetime of a battery is the result of the ageing processes ongoing within the battery. These ageing processes are irreversible changes in the components of a battery, in the materials used, or in the properties of the battery [12]. The ageing processes can be seen as being induced or furthered by stress factors, such as a battery remaining at low states of charge a long time, the charge factor, the time delay between full charges, or the operating temperature. Battery life prediction typically attempts to develop quantitative models capturing the relationship between stress factors and ageing processes. Operating conditions, user requirements, operating regimes, and battery design often lead to a combination of stress factors. Consequentially, multiple ageing processes must be considered simultaneously, in particular, as these multiple ageing processes may interact and interfere with each other. Because the effects of ageing involve the entire system and not just the battery cell, it is even more difficult to predict the lifetime of hybrid energy storage systems. Wenzl et al. categorized battery life models into models that reflect the impact of ageing on performance

* Corresponding author. Tel.: +1 573 578 8837; fax: +1 847 327 3896.
E-mail address: weigert@mst.edu (T. Weigert).

0378-7753/$ – see front matter
doi:10.1016/j.jpowsour.2010.10.075

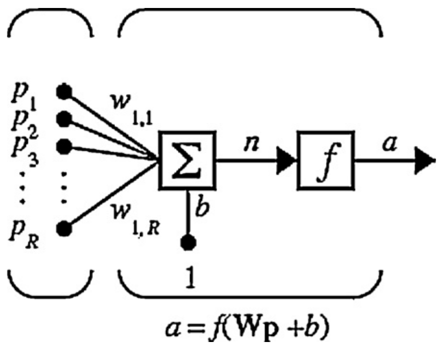

Fig. 1. Schematic depiction of a single neuron (processing element) of an ANN. The weighted inputs and the bias are summed (Wp + b), and a transfer function f is applied to the result, yielding the output of the neuron.

and models that link the lifetime to a measurable parameter [12]. However, many of the published models have not been validated experimentally.

In contrast, rather than deriving such models from physical relationships, estimation techniques attempt to infer a set of parameters that will minimize the sum-of-squared-differences between the observed behavior (target) and the model-generated output. Artificial neural networks (ANNs) are an effective estimation technique to approximate the parameters for complex, non-linear behaviors described by a combination of functions of the form $a = f(\mathrm{Wp} + b)$, where p is a vector of inputs, W is a vector of weights, b is a bias which is summed with the weighted inputs, and f is the transfer function applied. By combining such functions in series (linking the output of one function to the input of another), or in parallel (by applying an input to a set of such functions concurrently), or both, one can approximate arbitrarily complex non-linear functions. An ANN estimates the values of the weight parameters in these functions so that, given its input, the combined set of functions yields the desired output.

An ANN is composed of a set of processing elements, referred to as neurons. Each neuron corresponds to a function from one or more inputs to a single output, as described above (see Fig. 1). A number of neurons may be arranged in parallel into a layer, and such layers may be arranged sequentially, such that each output of a neuron of a given layer is connected as input to each neuron in the subsequent layers (for more complex systems there may also be back links between layers). The neurons of an ANN are therefore grouped into an input layer, zero or more hidden layers, and an output layer. Each input to the ANN is processed by a neuron in the input layer, each output is processed by a neuron in the output layer, and an arbitrary number of additional neurons may be inserted in hidden layers. There is no physical meaning to the neurons in the network, to the connections between the individual neurons, or to the arrangement of the network into layers. Similar to other estimation techniques, to avoid developing a system that works well for the training data but is not able to estimate new data, the simplest network that performs adequately should be constructed.

Through a training process, the weights and biases in the individual neurons are iteratively adjusted so that the overall network minimizes the sum-of-squared-differences between predicted outputs and observed outputs given a known set of training data. When training the ANN, an input vector and the observed output for that input vector are presented to the network, and the training algorithm adjusts the weights and biases of the network in the direction of the negative gradient, that is, in the direction where the mean squared error function decreases. A number of algorithms have been developed to perform this adjustment efficiently while avoiding being trapped in local minima.

Since ANNs can discover relationships between inputs and outputs of a system without a detailed understanding of the mechanisms involved, they can be effectively applied to estimate models for systems where such relationships are not clearly understood. Macdonald et al. studied the effect of temperature and running history (current, time step, time from start of test, ampere–hours, etc.) of a battery on its output voltage [13]. Gorman et al. utilized ANNs to simulate a single-load, constant discharge curve [14]. Grewal et al. investigated the effects of pulse current loads on the discharge time of a Li-ion battery and used a three-layer, feed-forward ANN to successfully simulate the state of discharge of the battery [15]. Kozlowski et al. predicted the remaining charge based on measurements (surface and internal temperatures, electrolyte pH, terminal and cell voltages, and electrical impedances) of the device obtained after an initial run. Impedance models provided additional information regarding the electrochemical processes occurring within the battery to the ANN [16]. Lee et al. used an ANN augmented with fuzzy reasoning and genetic algorithms to predict the state-of-charge of secondary batteries [17]. Farsi and Gobal used ANNs to predict supercapacitor performance, such as current densities, energy densities, and power densities, with good correspondence between predictions and the numerical model [18]. Parthiban et al. used ANNs to predict the shape of the curve of remaining capacity over 50 charge/discharge cycles under various experimental setups [19].

The lifetime of energy storage systems may refer to either the state-of-charge of the system (that is, how much usable capacity remains) or its state-of-life (that is, how many charge-discharge cycles remain). In this paper, we present an effective technique based on ANNs to predict the SOC of energy storage systems, both for batteries and for hybrid systems comprised of batteries and supercapacitors. By “lifetime” we therefore refer to a measure of time until the capacity of the system drops below a usable level in response to the application of a discharge pattern.

2. Experimental

The batteries used in these experiments were commercial alkaline primary batteries with a rated capacity of 1000 mAh. The supercapacitors (Taiyo Yuden, purchased from Digi-Key Co.) had a rated capacitance of 2 F, voltage of 2.3 V, and equivalent series resistance (ESR) of 50 mΩ. The ESR of batteries and supercapacitors was measured on a Solartron 1255 frequency response analyzer interfaced with an EG&G 263A. For the battery–supercapacitor hybrid system, each battery was connected to a supercapacitor in parallel. The batteries and hybrid systems were discharged on a Maccor battery test system. A high current pulse discharge was applied (from 1.5 A to 2.5 A). The pulse width of the current was fixed at 0.5 s. The duty ratio was set at either 0.17 or 0.25. The cut-off voltage was set at 1 V. We conducted three experiments, using different sets of primary batteries, with the pulse discharge load set at 1.5 A, 2 A, or 2.5 A, respectively. All tests were conducted under ambient conditions. In all experiments, we used both batteries and battery–supercapacitor hybrids together. We relied on primary batteries instead of secondary batteries in order to eliminate the impact of charging history and charging profile on lifetime.

We constructed an ANN to predict, based on inputs characterizing the device and its operating conditions, the number of pulse cycles before the voltage of the device drops below the cut-off voltage. We used a three-layer, feed-forward ANN comprised of an input layer with a neuron per input, a hidden layer, and a single neuron in the output layer which represents the predicted number of pulse cycles. The number of neurons in the hidden layer was selected based on the number of input neurons, following the rule of thumb to set the size of the hidden layer to roughly twice the size of the input layer.

A	Hybrid	Min	Max	Avg	StdDev
1.5	Yes	1084	3315	2104	529
	No	158	2241	1215	640
2.0	Yes	432	1451	755	240
	No	135	1152	561	245
2.5	Yes	88	673	288	120
	No	80	505	225	105

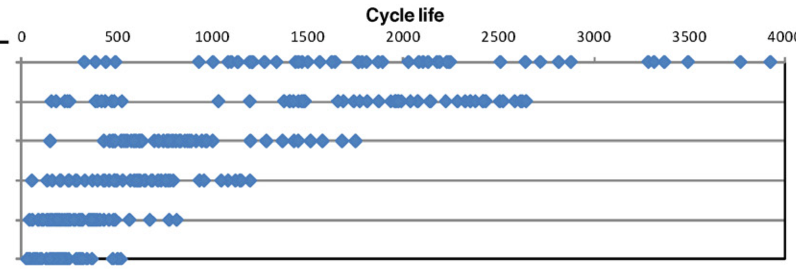

Fig. 2. Observed variation in lifetime: minimum, maximum, and average number of pulse cycles before cut-off voltage and standard deviation, for various devices (left) and graph of lifetime for various devices (right).

For training the ANN, we used the Levenberg–Marquardt backpropagation algorithm, which is an efficient implementation of a quasi-Newton method to estimate the network parameters. In order to improve generalization of the network and to avoid overtraining, we relied on the *early stopping* technique: A set of validation data is monitored during training. Initially, both the error in the training data and the error in the validation data decrease. When the network begins to overfit the data, the validation error increases. At this point, training is stopped.

In all our experiments, we divided the available data into three distinct data sets: training data, validation data, and test data. The network was trained with the training data, and the validation data was used to determine when to stop training. The trained network was then evaluated using the test data. Each experiment was comprised of a (large) number of runs, for various network configurations. As parameter estimation through ANNs is a stochastic process, in each run the network was trained and evaluated 100 times, and the average result was chosen. In this way, we avoided tainting the finally obtained results caused by accidentally good or bad performance on specific samples.

3. Results and discussion

A large variation in the measured lifetime for tested batteries and battery–supercapacitor hybrids, expressed as the number of pulse cycles before the voltage dropped below the cut-off voltage, was observed. Fig. 2 shows the range of minimum and maximum lifetime for the three experiments conducted at 1.5 A, 2.0 A, and 2.5 A, for batteries and hybrid systems, along with the mean and standard deviation for each experiment. For example, the first line shows that for the experiment involving a battery–supercapacitor hybrid with a current of 1.5 A applied, the observed life ranged from 1084 to 3315 pulse cycles, with a mean life of 2104 pulse cycles at a standard deviation of 529 pulse cycles. The graph on the right shows the distribution of the measured lifetime of the indicated devices. As Fig. 2 shows, a hybrid system increases the average life. However, the observed variation in lifetime is extremely large. ESR measurement on the batteries revealed a range from 200 mΩ to 600 mΩ, with an average of approximately 324 mΩ and a standard deviation of 69 mΩ. This wide spread in ESR may have a significant impact on the lifetime for high current drains.

In an initial experiment, we attempted to predict, from known characteristics of the devices, the lifetime in terms of the number of pulse cycles after which the voltage of the battery or battery–supercapacitor system drops below 1 V. For these experiments, we used ESR, duty ratio, and information as to whether the device was a battery or a hybrid as the input to the neural network. The neural network was configured with 3 input neurons, one for each of these characteristics, and a single output neuron indicating the predicted lifetime. The network performed best with 3 hidden neurons. Fig. 3 shows the performance of the ANN predicting the life of the systems based on known device characteristics only. As can be clearly seen, the performance of the network is rather poor: the slope of the correlation of predicted values vs. observed values is only 0.192, far from indicating a good match, and the correlation coefficient is below 0.5.

The lifetime of a battery is the result of the ageing processes ongoing within the battery. The aging processes can be linked to stress factors. Wenzl et al. [12] describe stress factors as statistical parameters calculated from a time series of operating conditions. These parameters link the operating conditions to the observed lifetime. For the subsequent experiments, we assumed that this linkage between operating condition and lifetime is expressed by the discharge curve. That is, rather than trying to infer the relationship between stress factors and lifetime, we attempted to infer the relationship between the shape of the discharge curve and the lifetime. If the discharge curve correctly summarizes this linkage, then we would be able to predict the lifetime of the battery. Such prediction is only useful if it can be made based on a relatively short initial segment of the overall discharge curve.

To predict the lifetime from the initial segment of the discharge curve, we configured the ANN with input neurons corresponding to the voltage measured at a given point of the discharge curve. We sampled the discharge curve every 5 pulse cycles. For example, to leverage the first 100 pulse cycles of the discharge curve, we set up the ANN with 20 input neurons, reflecting the voltage at the 5th, 10th, 15th, up to the 100th pulse cycle. We experimented with various configurations of hidden neurons, from 1.5 to 2.5 times the number of input neurons. The single output neuron yielded the

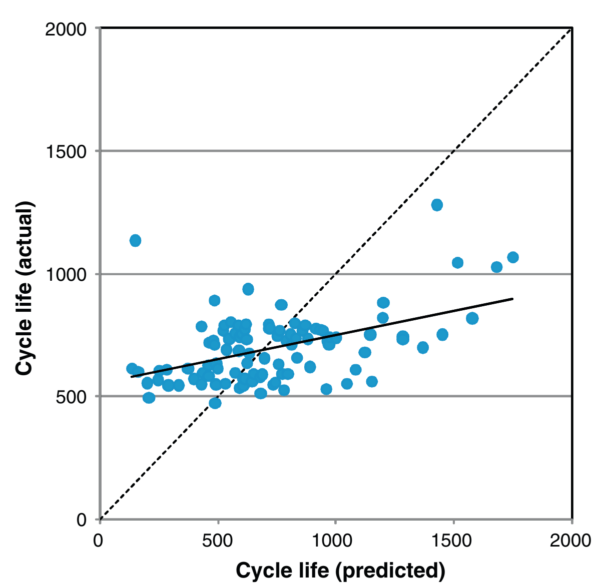

Fig. 3. Predicted lifetime based on readily measurable device characteristics (ESR, duty ratio, hybrid vs. battery) at 2.0 A pulse discharge. The horizontal axis shows the predicted pulse cycles; the vertical axis shows the actual pulse cycles.

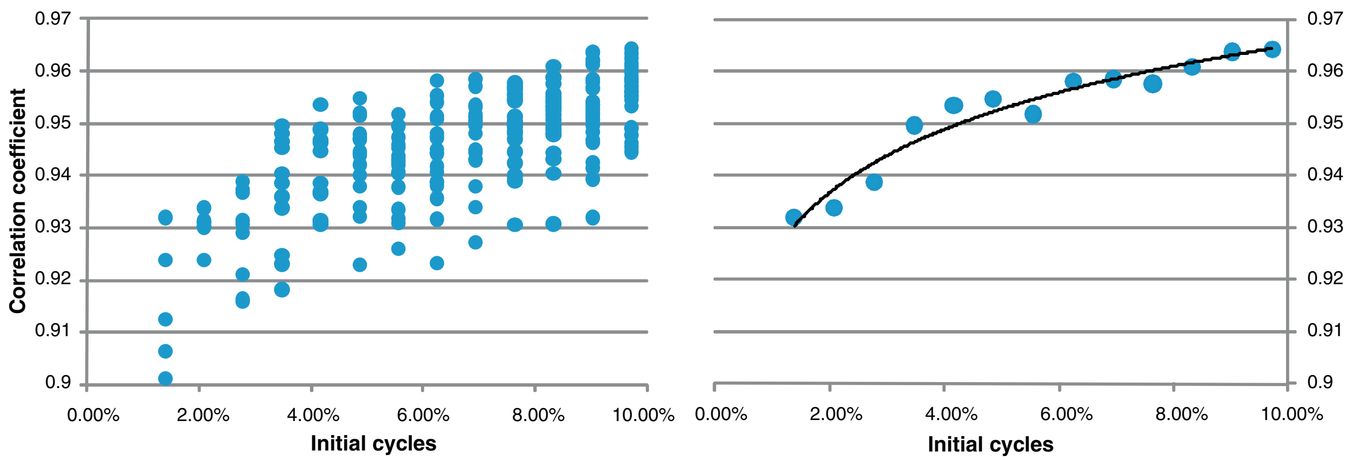

Fig. 4. Prediction results at 2.0 A pulse discharge using different lengths of data: (a) with varying number of hidden neurons and (b) best prediction. The horizontal axis shows the length of the initial segment of the discharge curve in percent of the average number of pulse cycles; the vertical axis shows the correlation coefficient.

prediction of the lifetime (that is, the number of discharge pulse cycles before the voltage dropped below the cut-off voltage).

Fig. 4a shows the predictions obtained for the experiment at a current of 2 A, when varying the number of hidden neurons and the initial length of the discharge curve. The horizontal axis shows the initial length of the discharge curve in terms of the percentage of the average life. The vertical axis shows the correlation coefficient of the obtained prediction. The different points for a given initial length of the discharge curve result from differently configured nets (using different numbers of hidden neurons). Not surprisingly, the prediction improved with the length of the observed discharge curve. As the length of the initial segment increased, the impact of the chosen configuration of the net became smaller (showing a smaller spread of the correlation coefficient). Fig. 4b shows only the best results for each chosen initial segment of the discharge curve. The predictions were quite good already at very short initial segments: at less than 2% of the average life, the prediction showed a correlation coefficient above 0.93. The accuracy improved rapidly and was above 0.95 after 4% of the average life. Above 8% of the average life, the correlation coefficient exceeded 0.96. In Fig. 5, we show the predictions at 4% of the average life: the fit between predicted and observed values is very good, and the slope of the correlation line is 0.857, sufficiently close to 1. We then compared the predictions obtained for the experiments at various levels of discharge current (at 1.5 A, 2.0 A, and 2.5 A; see Fig. 6). We see that the accuracy of the prediction increased with lower variation in the samples (see Fig. 2). For example, at 4% initial length the correlation coefficient obtained for the experiment at 2.5 A current, where the samples exhibited much less variation, was close to 0.98 (with a correlation slope of 0.916). In Fig. 7, predictions at 4% of the average life are shown, demonstrating excellent fit between prediction and observed result.

Our initial experiment had shown that we were not able to obtain good predictions from easily measurable device characteristics alone (see Fig. 3). We then conducted an experiment to determine whether knowledge of these characteristics could further improve the predictions based only on the initial segment of the discharge curve. We added input neurons to the network reflecting these characteristics. Fig. 8 shows the predictions obtained using 4% of the average life at a discharge current of 2.0 A when providing additional knowledge about various aspects of the device. As can be seen, the information about the ESR had little impact on the prediction accuracy. Adding information whether a device is a battery or a hybrid improved the prediction somewhat. Knowledge about the duty cycle did improve the prediction more substantially. Providing any additional information did not further improve the prediction. A similar improvement was obtained when providing information about the ESR jointly with whether

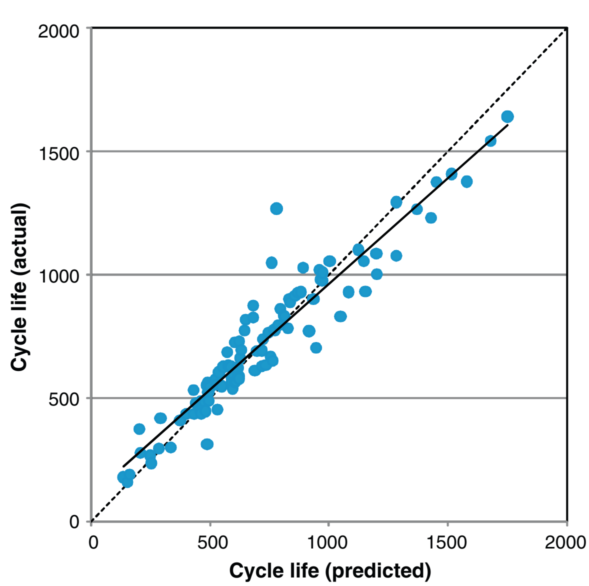

Fig. 5. Predicted lifetime based on initial segment of the discharge curve of roughly 4% of the average lifetime at 2.0 A pulse discharge. Axes are as in Fig. 3.

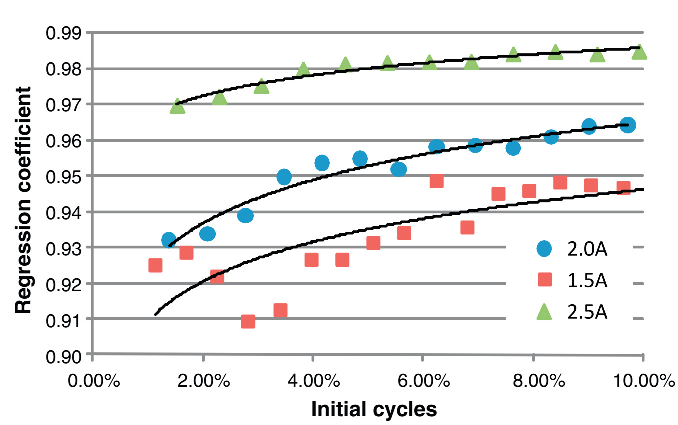

Fig. 6. Prediction results at 1.5 A, 2.0 A, and 2.5 A pulse discharge (curves from bottom to top). Axes are as in Fig. 4.

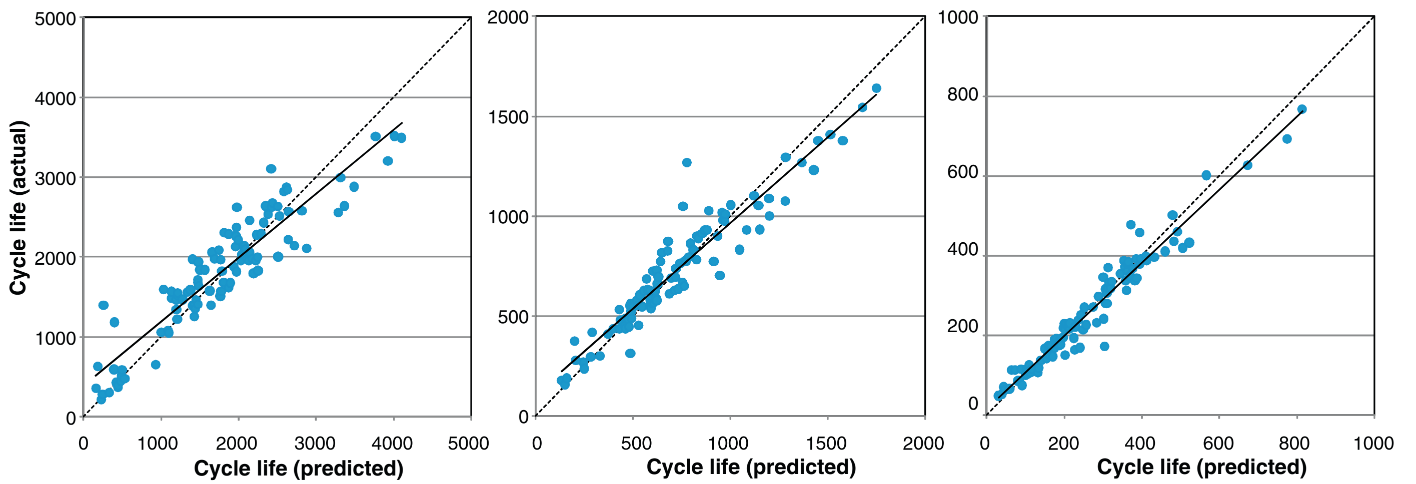

Fig. 7. Predicted lifetime based on initial segment of the discharge curve of roughly 4% of the average lifetime at 1.5 A, 2.0 A, and 2.5 A pulse discharge (graphs from left to right). Axes are as in Fig. 3.

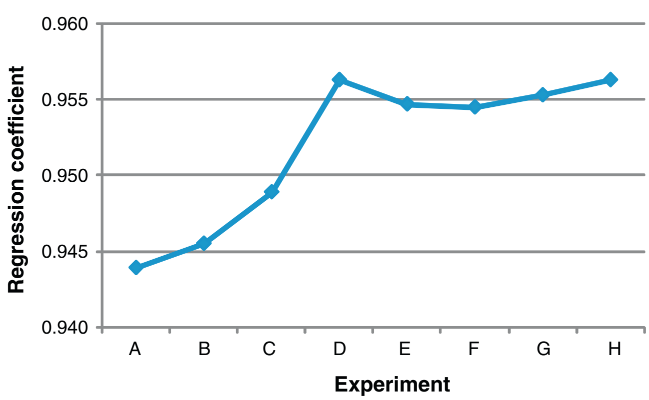

Fig. 8. Best prediction results at 2.0 A combining the initial segment of the discharge curve of 4% of the average lifetime and readily measured device characteristics for discharge data only (A), including ESR (B), battery vs. hybrid (C), duty cycle (D), battery vs. hybrid and duty cycle (E), battery vs. hybrid and ESR (F), ESR and duty cycle (G), battery vs. hybrid, ESR, and duty cycle (H).

the device was a battery or hybrid. We conclude that while most of the linkage between stress factors and lifetime is already reflected in the shape of the discharge curve, there are still aspects of this linkage that are reflected by device characteristics.

We were further interested in how resilient our method was to changes in the operating conditions or makeup of the device. In the experiment shown in Fig. 9a, we combined data from batteries, hybrid devices with 2 F capacitors, and hybrid devices with 4 F capacitors. We wanted to see whether predictions could be made by examining discharge curves from devices with different physical makeup. Excellent predictions where again obtained, as shown by the bottom curve. These predictions can be further improved by adding information as to which device was a hybrid and the capacity of the supercapacitor (see the top curve). The improvement of the results over the earlier experiment is probably due to having a much larger data set available to train the ANN and the network being able to distinguish the physical difference in the devices. In the experiment shown in Fig. 9b, we combined data obtained when discharging at 1.5 A, 2.0 A, and 2.5 A current, reflecting different usage characteristics. Again we obtained excellent predictions. The bottom curve shows the prediction based only on the initial segment of the discharge curve, the top curve shows the prediction when adding information as to which discharge rate was used. These experiments suggest that our method can accurately predict the lifetime even when the physical characteristics of the devices vary or when the usage behavior varies between samples.

Predicting the lifetime of batteries and hybrid systems is essential for designing and determining their applications. Predictive methods based on models derived from the internal chemistry of the devices have not yet been proven. Although it is impossible to follow the discharge profiles of most batteries and/or hybrid energy systems in real life, some important information about

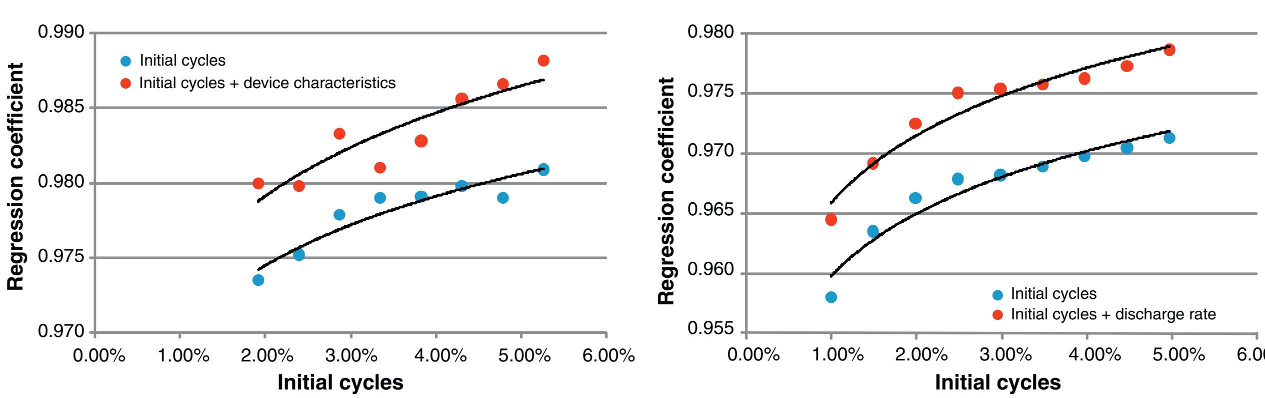

Fig. 9. Prediction results at 2.0 A from (a) batteries, hybrid devices with 2 F capacitors, and hybrid devices with 4 F capacitors combined into a single data set, and (b) at 1.5 A, 2.0 A, and 2.5 A pulse discharge combined into a single data set. Axes are as in Fig. 4. The bottom graphs show the prediction results based on the initial segment of the discharge curve, the top graphs show the prediction results based on the initial segment of the discharge curve together with information as to the characteristics of each device or mode of operation.

these systems can be extracted from accelerated life tests such as high current pulse discharge, which is a common practice for quality control with many battery manufacturers. However, this practice is time and energy consuming, especially for hybrid systems with even further extended lifetime. Therefore, the reported method of predicting lifetime based on a short initial segment of pulse cycle data is highly beneficial. To date, there are only a limited number of studies of the prediction of battery life, and even less for battery–supercapacitor hybrid systems. This work developed a method to predict the state-of-charge for energy storage systems that is applicable to both batteries alone and battery–supercapacitor hybrids. In follow-on work, we plan to extend this approach to state-of-life predictions using secondary batteries.

4. Conclusions

An approach to predict the lifetime of batteries and battery–supercapacitor hybrid systems using ANNs has been demonstrated. Our results show that, relying on only a small fraction of the discharge data (less than 4%), the ANN can predict the state-of-charge of these devices and systems with very good accuracy. The key information required to predict the lifetime is already captured in the discharge data itself, which reflects both the physical characteristics of the device as well as the conditions of operation (e.g., applied current, duty ratio, or whether the device was a battery or a hybrid system). ANNs are a powerful technique to effectively predict the life and discharge behavior of energy storage devices and/or systems without requiring a detailed investigation of the internal chemistries and interference between the chemistries of these devices.

Acknowledgments

We would like to acknowledge the financial support from Ontario Center of Excellence, Electrovaya Inc., and NSERC Canada.

References

[1] J. Zheng, T. Jow, M. Ding, IEEE Aerosp. Electron. Syst. Mag. 37 (2001) 288.
[2] G. Sikha, B. Popov, J. Power Sources 134 (2004) 130.
[3] J. Miller, Proc. Electrochem. Soc. Conf. 95–29 (1995) 246.
[4] L. Jarvis, T. Atwater, P. Cygan, J. Power Sources 79 (1999) 60.
[5] L. Jarvis, T. Atwater, P. Cygan, IEEE Aerosp. Electron. Syst. Mag. 15 (2000) 37.
[6] R. Dougal, S. Liu, R. White, IEEE Trans. Comp. 25 (2002) 120.
[7] C. Holland, J. Weidner, R. Dougal, R. White, J. Power Sources 109 (2002) 32.
[8] B. Conway, Electrochemical Supercapacitors: Scientific Fundamentals and Technological Applications, Kluwer Plenum, New York, 1999.
[9] R. Kotz, M. Carlen, Electrochim. Acta 45 (2000) 2483.
[10] T. Christen, M. Carlen, J. Power Sources 91 (2000) 210.
[11] A. Chu, P. Braatz, J. Power Sources 112 (2002) 236.
[12] H. Wenzl, I. Baring-Gould, R. Kaiser, B.Y. Liaw, P. Lundsager, J. Manwell, A. Ruddellg, V. Svobodah, J. Power Sources 144 (2005) 373.
[13] M. Urquidi-Macdonald, N.A. Bomberger, J. Power Sources 74 (1998) 87.
[14] C. O'Gorman, D. Ingersoll, R. Jungst, T. Paez, Proc. 31st Ann. Hawaii Intl. Conf. on System Sciences, vol. 5, 1998, p. 115.
[15] S. Grewal, D. Grant, Proc. Int. Telecommunications Energy Conf. (IEEE), vol. 484, 2001, p. 14.
[16] J. Kozlowski, M. Watson, C. Byington, A. Garga, T. Hay, Proc 55th Mtg. Soc. Machinery Failure Prevention Technology, 2001, p. 165.
[17] Y. Lee, T. Kuo, W. Wang, Proc. 35th IEEE Power Electronics Specialists Conf., 2004, p. 2759.
[18] H. Farsi, F. Gobal, Comput. Mater. Sci. 39 (2007) 678.
[19] T. Parthiban, R. Ravi, N. Kalaiselvi, Electrochim. Acta 53 (2007) 1877.