ELSEVIER

Journal of Power Sources 74 (1998) 87–98

JOURNAL OF POWER SOURCES

Predicting failure of secondary batteries

Mirna Urquidi-Macdonald *, Neil A. Bomberger

Engineering Science and Mechanics, The Pennsylvania State University, University Park, PA 16802, USA

Received 18 November 1997; accepted 13 January 1998

Abstract

The ability to predict the failure of secondary batteries is important. However, when determinism is not used to make the predictions because the complexity of the problem, difficult questions arise. Data analysts must always determine how much information is available in a given database and how much information can be squeezed from the database. A philosophical question is frequently posed: How long into the future can we predict based on past information? For the prediction of battery cycling life this question can be formulated as: How long must a battery (or cell) be tested to predict when it will fail? The answer to this type of question depends on how many variables define the problem, how much we know of the problem, how effective we are at squeezing information from the database, and how much knowledge and reliable data we have available to build a predictive model. The quality of the model will be measured by its ability to predict the future behavior of the system. The prediction of cycling life of batteries has been until now an impossible task. We are convinced that this is in part because the problem is very difficult, and in part, because the information available in databases has not been manipulated enough to produce a reliable predictive model. Models based on similar techniques are expected to have similar predictive capabilities. The methodology used in this project is now being used on an extensive database (thousand of hours) for NiCd batteries (NASA Goddard Space Flight Center data), and on a complete database (many variables are being controlled and measured) for Li/polymer batteries generated at the Battery Laboratory at Penn State.

Keywords: Life prediction of secondary batteries; Lithium–ion batteries; Artificial neural networks

1. Introduction

Some of the secondary batteries on the market (e.g., lithium-based) are promising because they offer a high-power density, and great mechanical flexibility. However, not all of the batteries built under the same specifications and with the same methodology will reach the expected long life. Until now, the reasons why these promising lithium–polymer technologies do not reach that desired long life ($\gg$ 1000 cycles) are not well understood nor do we know why the life of certain batteries belonging to more long-life family systems (lithium–ion) are shortened while other batteries of the same family last a long time.

It would be of great interest especially for low-weight applications to be able to predict the performance life of a battery by running a short series of tests on it. One such application is space flight. It is important to minimize the weight of the launch vehicle. For these applications, tests are run on many cells (one or several cells may form a battery system), and statistics are used to calculate how many cells will be included in a flight, using information on the number of cells needed to supply the required power, the probability that cells will fail, and a safety factor of extra cells. These extra cells are excess weight for the launch. If a method could be developed to determine whether cells will fail by running a short series of tests on individual batteries, the weight and volume of the battery pack could be reduced.

It is our intention in this paper to present and discuss a methodology aimed at squeezing information from a database. The model learned from a small database containing 2 h of cycling life of lithium–ion secondary batter-

* Corresponding address. mumesm@engr.psu.edu.

0378-7753/98/$19.00

ies. The model used short-term information (60 and 10 min) to predict longer term information (2 h) for which measured data exist; that actual data was used to assess the predictive performance of the model. The model was then used to explore the system's longer-term response (10 h) for which measured data was not available. Extrapolation of knowledge (expectations based on experience) was used to assess the system performance.

Two types of current/time patterns were imposed on a battery (10 cells) and voltage/time response patterns were measured. The patterns were imposed at three different temperatures. This research effort identified those variables collected during life testing of the battery system that contributed the most to predicting the performance of the battery in the future. The prediction model used a pattern recognition artificial neural network (ANN). Data were presented to the net as inputs (current pattern) and outputs (voltage pattern) during the net training. The way the data are represented is very important and defines the net's learning speed. In general, we can state that to squeeze information from the data, it is very important to represent the data in various ways. For example, we learn music by reading it but also by expressing it through a musical instrument (transformation). Data transformation and data representation are key issues in the learning process and in squeezing information from the data. In this paper, we discuss the advantages and disadvantages of three different data representations.

The objective of this paper is to use artificial neural networks, specifically feed-forward backpropagation (FBP) networks, to explore the possibility of predicting cycling life failure; based only on short-term data. The research work is based on certain predictive capabilities shown by artificial neural networks [1].

The predictive capabilities of FBP nets are presented and discussed. Three data representations of the inputs and outputs used to train the network are discussed: an amplitude-time representation (data representation one), a frequency representation using a discrete Fourier transform (data representation two), and a wavelet coefficient representation using wavelets to compress the current and voltage signals (data representation three). In this paper we also present an analysis that aims to categorize the variables measured during battery testing in order to determine their relative importance in predicting the cell cycling life.

2. Neural networks: An overview

Artificial neural networks are a mathematical tool with excellent capabilities for pattern matching, recognition, and classification. ANNs are inspired by the biological behavior of neurons. The fundamental building blocks of biological neural (nervous) systems are cells called neurons. The neurons are represented by a summation and a transfer function. An artificial neural network may also consist of layers of interconnected neurons. A network is composed of one output layer, one input layer, and one to three hidden layers. Signals travel from the input layer, through the hidden layers, to the output layer. Only two hidden layers are usually needed to learn non-linear transformations. Each hidden layer in a network adds one degree of non-linearity.

There are two main types of nets: supervised and non-supervised. Supervised nets learn by examples, and during the training process the weights or connections in-between neurons change with time (the transfer functions remain practically constant) so that the net as a whole adapts to reach an answer close to the desired output (presented as a single value, vector, or matrix of information) given an input vector. Once the net learns 'a job' the weights remain constant. During the learning process, the net establishes relationships (highly non-linear) between the inputs and the outputs; once the net is trained, that categorization of the inputs by weight-importance or 'contribution to the learning process' can be used to explore the importance of the input variables in reproducing the desired output.

For this research a supervised type of learning was chosen; the input and output data were run through the net during training. The learning was accelerated with additional information in the form of input series calculated from the input data (functional links). In this work we explored several types of supervised net designs; the results of only one of them are reported.

Feed-forward backpropagation (FBP) networks are composed of multiple layers of processing elements (PEs) that are represented by transfer functions such as sigmoids. The network consists of an input layer, and output layer and any number of hidden layers that give the network the ability to solve non-linear problems. No more than two hidden layers are usually necessary to learn a difficult task. The input layer sends all of the input values through weighted connections to the first hidden layer, where these inputs are summed and the processing elements use the transfer function to produce an output. This output is sent through weighted connections to the next hidden layer, where the PEs similarly produce an output, which is finally sent to the output layer. The network undergoes supervised training, which means that the desired output has to be known for each given input, and the network adjusts its weights to produce an output close to the desired output. The FBP network uses an error gradient descent learning rule to properly adjust its weights. The error between the network output and the desired output is expressed in terms of the network weights, so the gradient of the error with respect to the weights can be found. The weights change by descending this error surface along the gradient.

For this project, a commercially available neural network building package was used to design and create the

FBP networks used. The hyperbolic tangent (tanH) was used as the transfer function. The Extended Delta–Bar-Delta (DBD) net was used. The DBD net adds a momentum term to the error gradient descent learning rule to avoid the net response getting trapped in local minima of the error surface. The networks had two hidden layers: the first had twice as many elements as the input layer, and the second had the same number of elements as the input layer. The number of elements in the input and output layers was determined by the size of the input and output vectors in the data representation being used.

3. Lithium–ion battery tests

Battery test data were received from tests run on lithium–ion secondary batteries by SAFT America under contract to NASA Goddard Space Flight Center. These tests were conducted to evaluate lithium–ion secondary batteries as a power source for the pistol grip tool (PGT). This is a tool for tightening and loosening fasteners when repairing satellites in orbit, such as the Hubble Space Telescope. The PGT requires a minimum voltage of 28 V to operate, so a drop below this voltage would be considered a failure. A series of tests were run on a 10-cell lithium–ion battery using two different current discharge profiles and three different temperatures. The two current discharge profiles were a typical current discharge profile (Fig. 1—profile 1), and a worst-case discharge profile (Fig. 2—profile 2). The typical current discharge profile has current spikes ranging from 2 to 8 A in groups of four to simulate the high torque needed to initially loosen a fastener, and longer current discharges at 1 and 1.5 A to simulate the removal or tightening of a fastener. The worst-case current discharge profile has current spikes ranging from 4 to 10 A and an extended current discharge plateau at 2.5 A to simulate the deployment of a solar array. These discharge profiles were applied to the battery in the order profile 1, then profile 2, at temperatures of 20, 0, and 40°C, and the running time, current, output voltage, and internal temperature of the battery were measured. Between each test, the battery was recharged to full capacity. These tests, including the recharge cycles, were labeled alphabetically. The data received for the discharged cycles was from test B (profile 1, 20°C), test D (profile 2, 20°C), test H (profile 2, 0°C), and test J (profile 1, 40°C).

4. Data representation one

The format in which the data are presented to an FBP network is important in determining whether the network learns effectively. In this case, the data contained a record of the current, running time, temperature, and resulting voltage output of the battery at discrete moments in time. Merely presenting a datum point of the current, running time, and temperature of the battery to the network would not be sufficient for the network to learn the relationship between these few inputs and the output voltage. Among variables that impact cell behavior other than the instantaneous current and temperature are the previous currents and the previous 'over a threshold' current value that the cells experience. Those variables affect the cell performance at any given moment. In order to train the net effectively, information about these variables also had to be presented to the network.

Four sets of data from the tests (A, D, H, and J) performed on secondary lithium–ion batteries were used to train a backpropagation network. For each battery test, the current, time step, running time, ampere-hours, temperature, and resulting voltage output for a 10-cell battery stack were recorded.

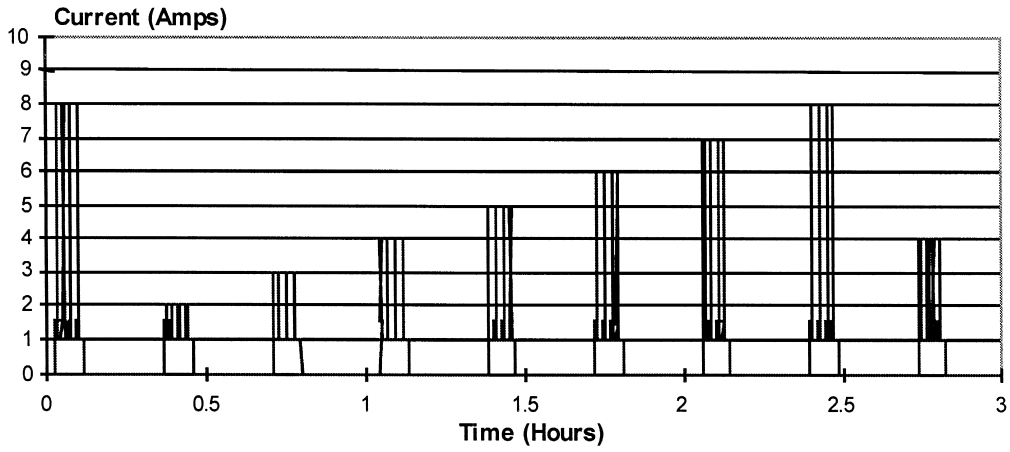

Fig. 1. Discharge profile 1: Typical current discharge profile. Test B at 20°C, and test J at 40°C.

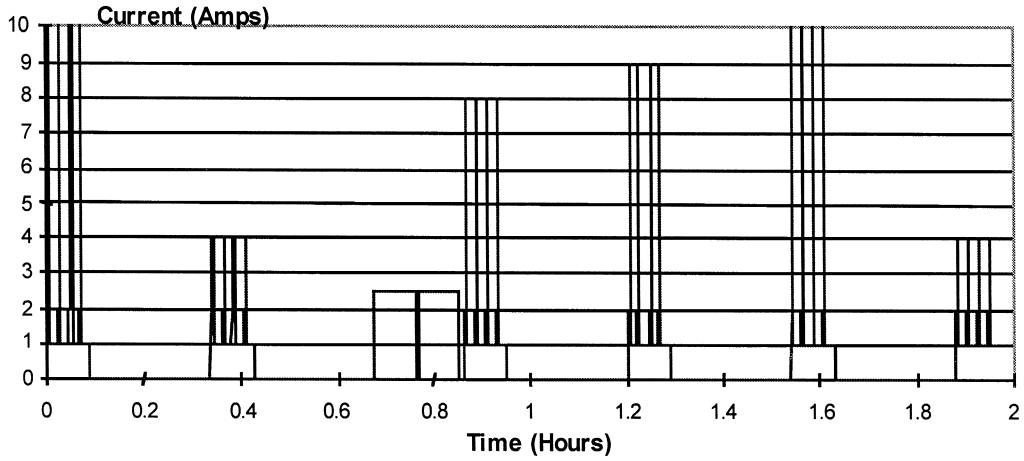

Fig. 2. Discharge profile 2: Worst-case current discharge profile. Test D at 20°C, and test H at 0°C.

Data representation one contains the following inputs: current (I), time step (Δt), time from start of the test (t), ΔI, $|\Delta I|$, ampere-hours, ampere-hours (resets at 0 A), $\Sigma(\Delta I)$, $\Sigma|\Delta I|$, temperature (T), $\log(T)$, a counter of the current peaks in the ranges 4–6 A, 6–8 A, and 8–10 A, and a current history vector recording the time spent in the current ranges 0–1, 1–4, 4–8, and 8–10 A. Functional links of the form x^n and $x^{1/n}$ (where x represents any of the variables described above and $n = 2, 3, 4$) were also included. These functional links add information to the network and help it to obtain more efficiently the relationship between inputs and outputs [2].

The composite training file contained 54 inputs and 1 output. Each input vector corresponded to a single current value applied to the battery and a single voltage response measured (1416 data points were measured in a 2-h test). Sections of these files (also called tests B, D, H, and J) were used to train a backpropagation network with a first hidden layer of 108 processing elements and a second hidden layer of 54 processing elements.

5. Results of data representation one

The network was trained using several representations of current and temperature as described above, and voltage was the output. We had arbitrarily chosen to train the net using the first 60 min of tests B and D, and 10 min of tests H and J. Tests B and D consist of tests run with discharge profiles 1 and 2, respectively, both at 20°C. Tests H and J consist of tests run with discharge profiles 2 and 1, at 0 and 40°C, respectively. The network had to be trained with both profiles (tests B and D) so that the network could learn the relationship between the different current discharge profiles and the output voltage; and the temperature dependence from tests B, H, and J. 10 min time patterns

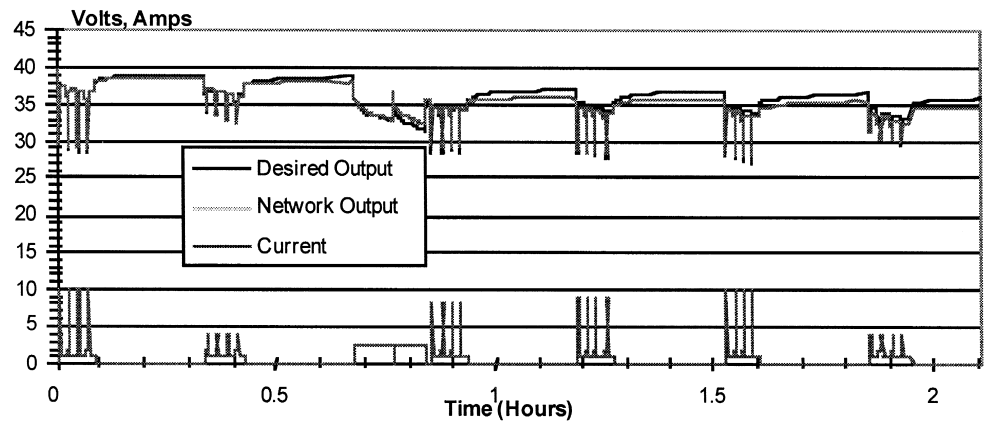

Fig. 3. Net voltage response predictions on 2 h of test H (light gray) compared to measured cell voltage response (black). The input current pattern is shown at the bottom of the graphic (dark gray, bottom). The network was trained with the first 60 min of tests B and D, and first 10 min of tests H and J (data representation one).

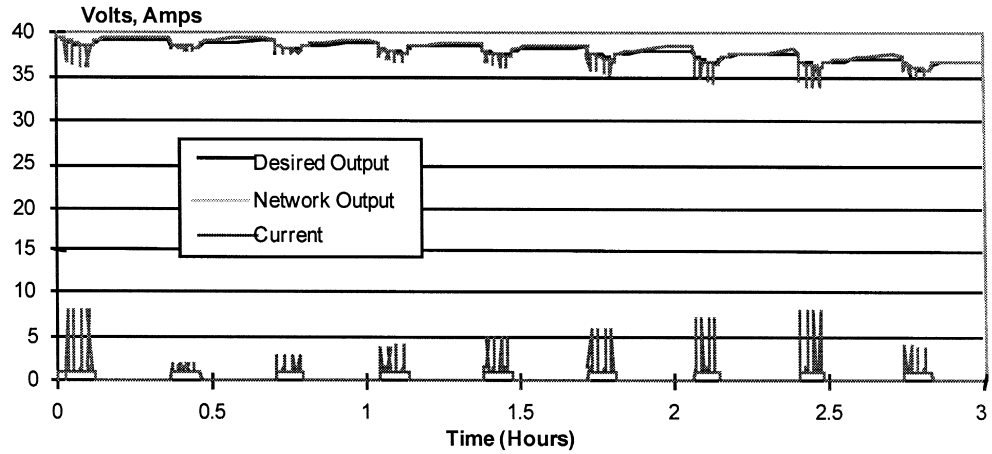

Fig. 4. Net voltage response predictions on 2 h of test J (light gray) compared to measured cell voltage response (black). The input current pattern is shown at the bottom of the graphic (dark gray, bottom). The network was trained with the first 60 min of tests B and D, and the first 10 min of tests H and J (data representation one).

were necessary for the net to learn the dependency in profiles 1 and 2, because the 2.5 A current plateau was imposed on the battery (Fig. 2) in the first hour.

This trained network was tested using the full test H (2 h) results shown in Fig. 3. The results for the network tested on full test J (2 h) are shown in Fig. 4. The net predictions were compared to the measured voltage response. In both cases, the network output voltage conformed very closely to the actual output voltage.

The net had to be trained with portions of tests H and J in order to learn the temperature (other than 20°C) dependence of the voltage output. We also tried training the net using the whole 120 min of only test B (profile 1 and 20°C); we then test the net by using the 120 min of test H (profile 2, at 0°C). The results are shown in Fig. 5. The net that was trained with this representation using only one temperature (20°C) and one applied current profile (profile 1) did not learn to differentiate well between different input profiles nor different temperatures. When the net was tested with a different profile and temperature, it had problems identifying the proper response. The net did not perform well in describing the response of the cells when a constant 2.5 A current was applied over a longer period of time. The net did not predict the correct response (deep voltage spikes) to the deep current spikes corresponding to the cells' response at low temperature, but instead the net prediction response was more modest (corresponding to what the net learned at 20°C).

To perform well using this representation, the net must be trained using 'best' and 'worst' profiles and 'best' and 'worst' environmental conditions. The net has no problem interpolating a given response, but it underestimates the extrapolated cell behavior. The results appear to indicate that under this representation, to extrapolate the cell behav-

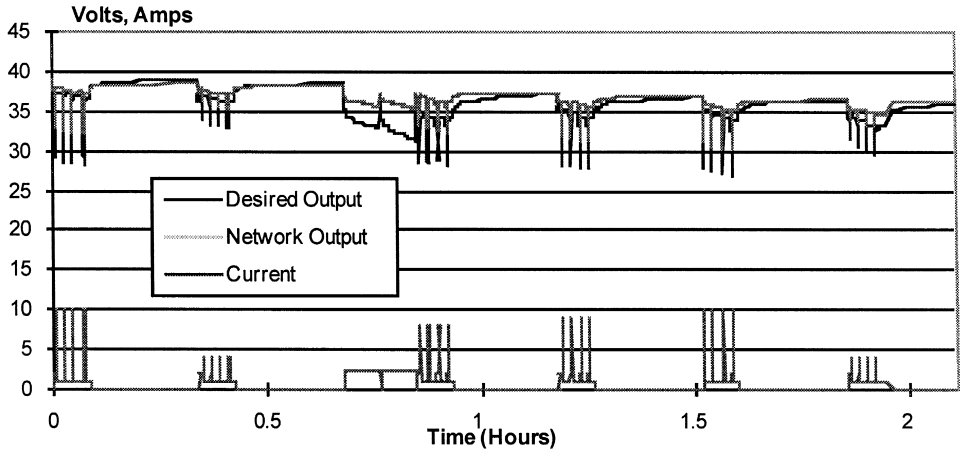

Fig. 5. Net voltage response predictions on 2 h of test H at 0°C, and profile 2-(light gray) compared to measured cell voltage response (black). The input current pattern is shown at the bottom of the graphic (dark gray, bottom). The network was trained with the 120 min of tests test B (data representation one and 20°C).

ior properly, it is necessary to have some measured information about the expected response of the cell to parameters that impact the cell cycling life (such as temperature and type of input current pattern applied to the cells). If the inputs used to train the net change from an example to the other and do not remain constants, the net will learn the trends and to extrapolate behaviors properly.

A network was also trained with all the data available from all four tests, then tested with extended time input patterns to get an idea of the long-term voltage performance of the cell. The hypothetical current input pattern used for training the net was constructed by placing input of tests B, D, H, and J in series and making a cumulative, continuous input, which represented a simulated 10-h battery test without recharging. The net was tested in a hypothetical pattern similar to test B (pattern 1, 'typical case'—the testing pattern was BBBB), but using a temperature sequence of 20°C for 5 h, followed by 2 h at 0°C, and finally followed by 3 h at 40°C. The net voltage response was expected to have a lower voltage rate at high temperature than at low temperatures (experimentally observed, but not shown here [3]). The depth of the spikes of the voltage signature was expected to be smaller at high temperatures than at low temperatures. Fig. 6 shows the input pattern applied to the trained net and the output voltage pattern predicted by the net. The net indeed produced lower output voltage spikes at the end of the simulated test corresponding to a temperature of 40°C than the section (between 5 to 7 h) corresponding to 0°C. The slope of the voltage-time response (or voltage rate) is more negative for low temperatures (0°C) than for higher temperatures (40°C), as expected.

One of the disadvantages of data representation one is the time required to pre-process the data before training the net. To avoid this problem, we propose to use as input to the net a compressed representation of the whole signature (2 hour). The problem that we may face in this new approach is that few hours of data may use a large number of data points. The following two data representations consider the use of the 2 h whole signature as input and the whole voltage signature as output. An identically-sized signature must be used to train the net and to test it; the net produces an output voltage response of the same size as the one used during the net training process.

6. Data representation two

Data representation two was developed by considering the discrete Fourier transform of the current and voltage patterns. In this representation, the entire signature of the current and voltage patterns of a test were transformed to a series of Fourier coefficients, which contain a smaller number of Fourier coefficients than the number of data points used to represent the signature. The Fourier coefficients of the current and voltage (as inputs and outputs, respectively) are presented to the network during training.

To reduce the training time of the network we sought to use only a few coefficients. Accordingly, the first 16 Fourier coefficients representing the signature (current or voltage) were used, resulting in 32 inputs and 32 outputs (the coefficients have real and imaginary parts). One problem with using only 16 coefficients was that the first 16 coefficients only roughly approximated the shape of the current and voltage patterns. For this reason, the network would not be able to duplicate the detailed measures of current or voltage. Because we may be interested in changing the time length of the input signatures, the time of the test (or time-length of the signature) was used along the

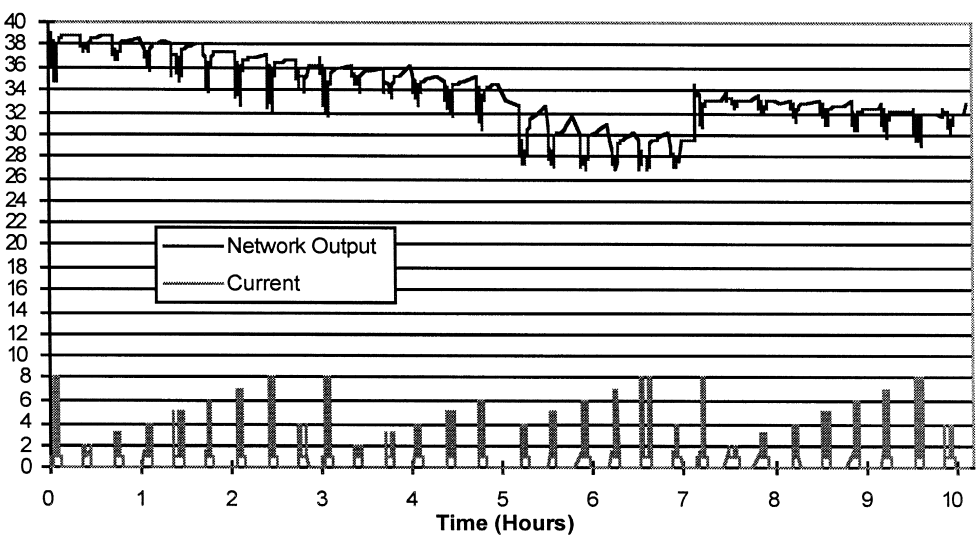

Fig. 6. Simulated network voltage response (black) for an input excitation pattern as indicated (gray). The net was trained with tests B, D, H, and J corresponding to profile characteristics 1 and 2 and temperatures 0, 20, and 40°C. The results of the net were obtained by inputting a current pattern as test B but for which the temperature was artificially assigned as 20°C for the first 5 h, 0°C for the 2 following hours and finally 3 h at 40°C.

temperature as inputs, bringing the total number of inputs to 34.

The advantage of this representation over the previous one (representation one) is that pre-processing the input data was fast and easy. A disadvantage was that for the current and voltage profiles to be represented exactly, we had to use as many Fourier coefficients as sampled test datum points (1416 for 2h). Another disadvantage was that in addition to recording the size of the input signature (current), the network should be trained and tested with Fourier coefficients from samples of the same length. This way, a Fourier coefficient input or output element will represent the same frequency for every vector that is presented to the network. If this is not the case, an input neuron (or process element, PE) will receive Fourier coefficients representing different frequencies, and the network may not learn the proper relationship between input and output frequencies.

The equation used to obtain the first 16 Fourier coefficients was

$$a(k) = \frac{1}{N} \sum_{n=0}^{N-1} u(n) e^{\frac{-j2\pi kn}{N}} \quad (1)$$

where N is the number of sampled points in the pattern, j is $\sqrt{-1}$ $u(n)$ is the value of the current or voltage at a sampled point, and $0 \leq k \leq 15$. The equation used to reconstruct a pattern from the coefficients $a(k)$ was

$$u(n) = \sum_{k=0}^{15} a(k) e^{\frac{j2\pi kn}{N}} \quad (2)$$

where the reconstructed current or voltage pattern consists of the real part of $u(n)$. A MatLab program was written to find the coefficients, and another MatLab program was written to reconstruct the pattern from the coefficients.

7. Results of data representation two

On the analysis that follows, one has to remember that a smaller number of Fourier coefficients will not nearly capture the detailed information on large databases.

To ensure that the Fourier coefficient inputs represented the same frequency in each input vector, the first 2 h (the length of the shortest test) of each test were used to find the Fourier coefficients. Then, a network was trained with tests B (profile 1, 20°C). The trained net (Fig. 7) was tested using the current pattern of test D (profile 2, 20°C). The net was trained on the 'typical case' current profile and tested on a 'never seen' input pattern from the 'worst case' current profile test. The 16 Fourier coefficient representation was a very poor representation of the real imposed current pattern, the Fourier representation lost a great deal of detail. Although the net predictions were very close to the Fourier representation of the measured voltage, the prediction does not have the level of detail required to determine whether the net 'sensed' the impact of the cell when the worst case current pattern was imposed on the battery instead of the typical case.

We also trained the net by using the four available files (B, D, H, and J) over a period of 2 h each. The current and voltage pattern corresponding to each of the files was represented by 16 Fourier coefficients. Once the net was trained, and could reproduce the Fourier representation of the measured voltage pattern signature very closely, we looked for the weight that had the highest absolute values, indicating the greatest contribution to the learning. By averaging the effect of changes in each input (the Fourier coefficients), the most important inputs—the inputs that

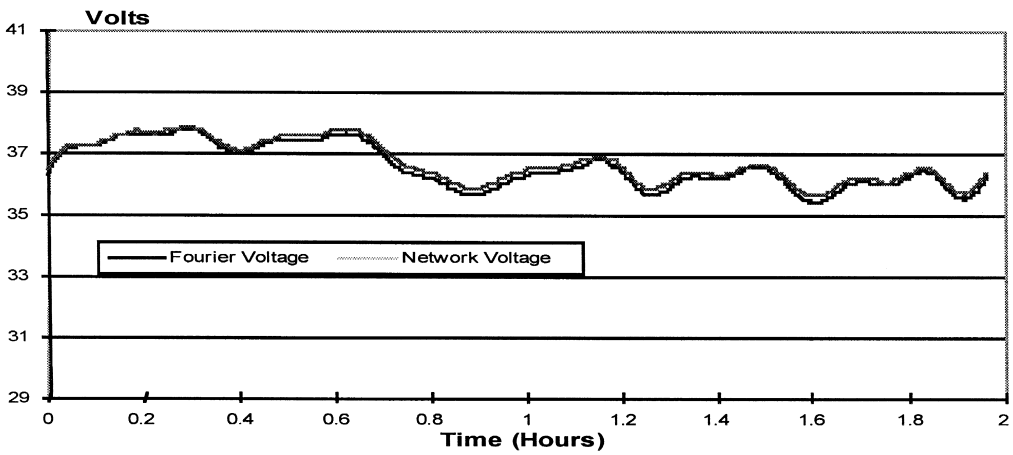

Fig. 7. Fourier voltage (black) and network voltage (gray) during a 2-h test. The network was trained with Fourier coefficients (16) on the first 2 h of test B (profile 1, 20°C); and tested with test D (profile 2, 20°C)

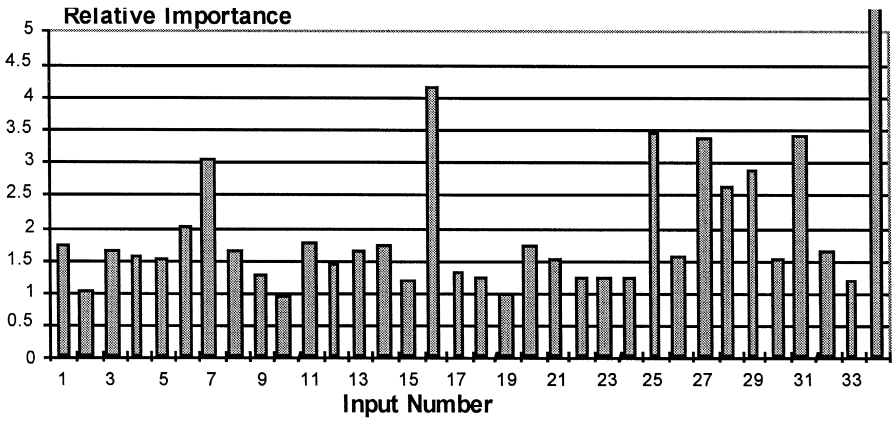

Fig. 8. Relative importance of inputs for network trained with Fourier coefficients (16), test duration, and temperature. The net used during training was the first 2 h of tests B, D, H, and J.

contributed the most to the learning—were found. Fig. 8 shows the relative importance for the inputs used to train this network. Inputs 1 through 16 correspond to the real current Fourier coefficients, 17 through 32 to the imaginary Fourier coefficients, 33 to sample time length (in our example, 2 h), and 34 to temperature (in our example 0, 20, or 40°C). Temperature is by far the most important input. This result consisted what we anticipated, because the only difference in the inputs between tests B and J and tests D and H is the temperature, so it is the sole cause of the difference in voltage output. Since the tests were conducted in series with recharge cycles between them, the resulting capacity loss due to discharge/recharge cycles would also have an impact on the voltage output. However, this effect is expected to be small over 2 h and is difficult to estimate from the data since the tests were conducted at different temperatures, so no information about this capacity loss was included in the input data. The next most important inputs were the later Fourier coefficients, which have highest frequencies. This is also logical because the coefficients with higher frequency correspond to the sharpest changes in the current patterns and had the most effect on the voltage output.

The network trained under data representation two was efficient in approximating the general shape of the Fourier voltage representation; however, the fine details were lost when the data is transformed using only 16 Fourier coefficients.

8. Data representation three

Data representation three was developed using the wavelet method of data compression. Wavelets represent an efficient way to compress data without losing too much information. Wavelets were used to compress 1416 data points representing the imposed current patterns and measured voltage patterns. The 1416 data points represented 2 h of real time testing on the battery system under different conditions (two applied current signatures, and three temperatures), as indicated in tests B, D, H, and J. These 1416 data points were compressed to 106 data points for each test by using a double compression scheme of a sparse representation (1/3) of the total signal. The wavelet used was a discrete biorthogonal exhibiting the property of linear phase, which is needed for signal and image reconstruction [4]. The first compression yielded a high frequency signal and a low frequency signal. The low frequency signal contained more information. Accordingly, the low frequency signal was compressed by forming a sparse vector (1 of 3 data points were retained); the resulting signal was compressed again resulting in 106 coefficients that were used to represent the original 1416 data points for each of the current or voltage signals from the four tests.

The signal reconstructed by using the wavelet captured most of the main features of the original signature (current/time or voltage/time) and the same scale by only using 106 data points instead of 1416, as can be seen for test D in Fig. 9. Data compression is an important step to perform before designing a neural network. The number of neurons in a layer of a neural network depends on the number of inputs and outputs; accordingly, it is more efficient to work with an input vector that does not have many elements (network designs containing more than 200 hundred neurons per layer become time-inefficient to operate on a PC). A compressed representation of the 2-h time-length vector reduces considerably the size of the neural network and, accordingly, the training time.

When compared to a Fourier transform with the same number of coefficients, wavelet representation gives superior results. The Fourier representation of the pattern (current or voltage) obtained does not come close to capturing the information that a wavelet representation captures.

9. Results for data representation three

A network was trained with 106 wavelet coefficients obtained by compressing the current and voltage signal or

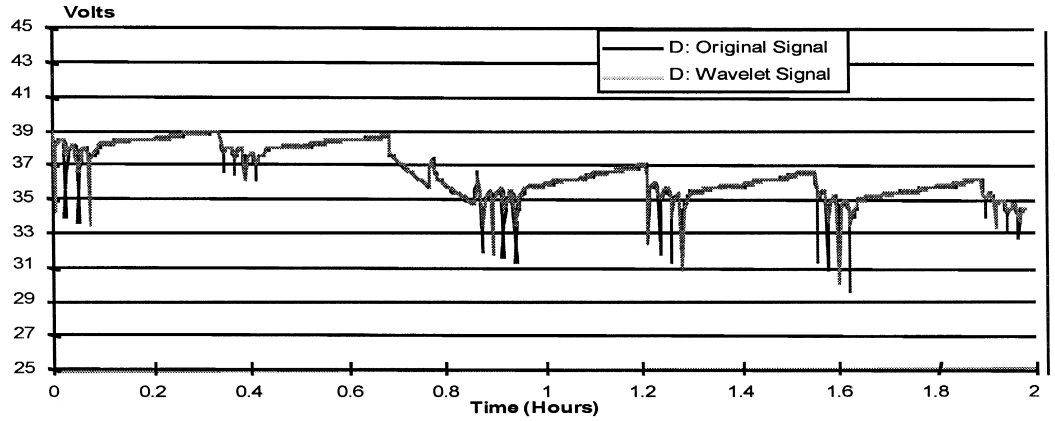

Fig. 9. Original voltage signal from test D (black) compared to voltage signal produced by wavelet compression using 106 coefficients (gray), representation (106 coefficients) of test J (profile 1 at 40°C).

pattern (2-h time-length) of test B. Then the network was tested with the 106 wavelet coefficients of current pattern from tests D, H, and J. Since these tests all occurred for different lengths of time, they were all trimmed to their first 2 h, so that respective wavelet coefficients for each test would refer to the same sections of the test.

Figs. 10–12 show the results for the network tested with tests J, D, and H, respectively. The initial and final data points of the file produced severe oscillations on the wavelet representation because the information was truncated at those two ends; the results of the net for those end points should be ignored.

Ignoring those end points, Fig. 10 compares the network prediction (gray) of the wavelet representation to the wavelet representation of the measured voltage (black). The training of the net was performed at 20°C under a current profile 'typical case' and tested on a current profile 'typical case' but at 40°C. The network predictions followed the voltage response of the cells to the input current profile very well. This input current profile was characterized by small spikes when the measures were performed at high temperature. The training contained the wavelet representation of the current and voltage signatures measured at 20°C under 'typical case' current profile. Therefore, the only way to explain the capability of the net to predict the modest response of the cells when the same profile of current was applied but at high temperature (as expected) is that the net learned the correct function that describes the temperature effect on the voltage response.

Fig. 11 compares the network prediction (gray) of the compressed wavelet representation and the wavelet representation of the measured voltage (black). The training of the net was performed at 20°C under a current profile 'typical case' and tested on a current profile 'worst case'

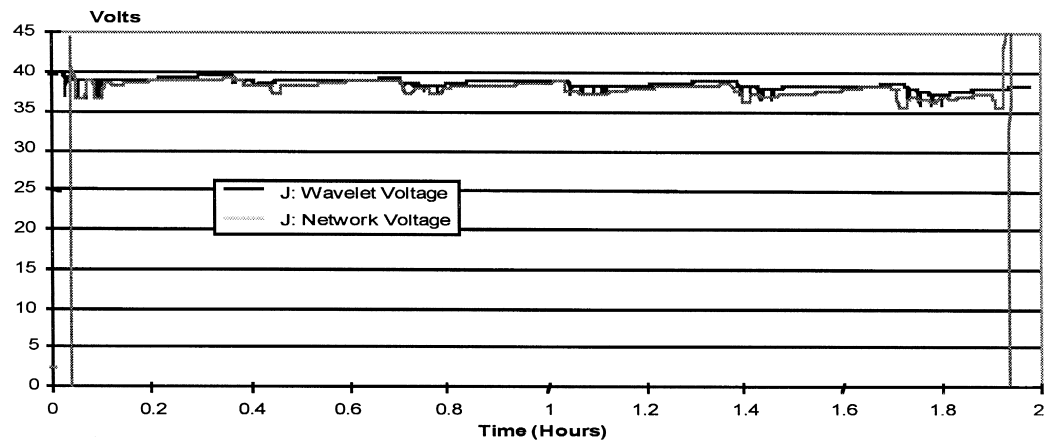

Fig. 10. Voltage response network predictions (gray) compared to the measured voltage response (black). The net was trained using test B (profile 1 at 20°C) with only 106 coefficients extracted using a wavelet representation of the original signal. The net was tested using the wavelet representation (106 coefficients) of test J (profile 1 at 40°C).

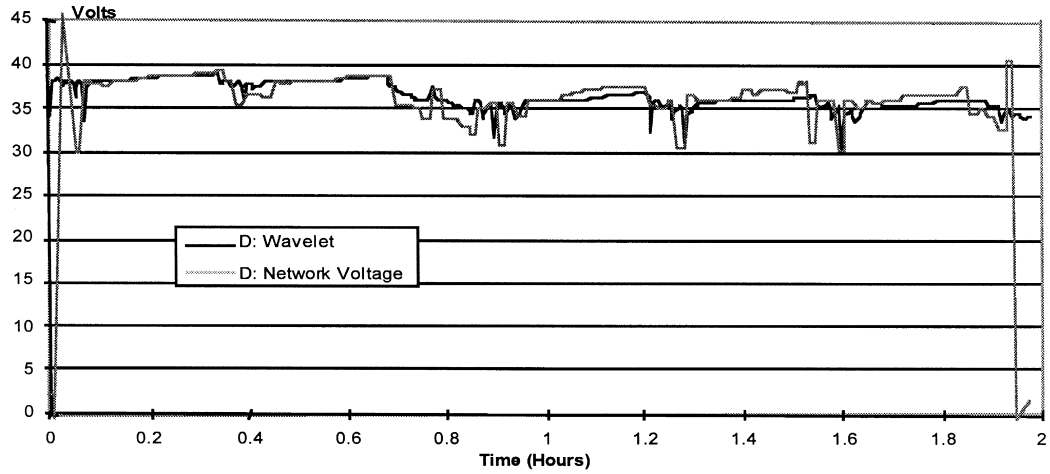

Fig. 11. Voltage response network predictions (gray) compared to the measured voltage response (black). The net was trained using test B (profile 1 at 20°C) with only 106 coefficients extracted using a wavelet representation of the original signal. The net was tested using the wavelet representation (106 coefficients) of test D (profile 2 at 20°C).

also at 20°C. The network predictions closely followed the voltage negative slopes (e.g., 0.7 to 0.9 h, see also Fig. 2) corresponding to the response of the cells to the 2.5 A plateau current profile; the net overestimated those plateaus, but it detected and differentiated them from the regular spikes very easily. The net training contained the wavelet representation of the current and voltage signatures measured at 20°C under a 'typical case' current profile. Therefore, we explained the capability of the net to predict the negative slope-plateau voltage responses occurring at times when the current plateaus were imposed (2.5 A imposed between 0.7 and 0.9 h time sample) by assuring that the net learned the correct function describing the voltage response of the cell.

Fig. 12 compares the network prediction (gray) of the wavelet representation and the wavelet representation of the measured voltage (black) when the net was trained at 20°C under a current profile 'typical case' and tested on a current profile 'worst case' at 0°C (lower temperatures cause cells to respond to current spikes by producing larger voltage spikes than that obtained at high temperatures). The network closely predicted the large spikes at 0°C (compare Fig. 12 to Fig. 10); the spikes were more pronounced and closely follows the measured voltage re-

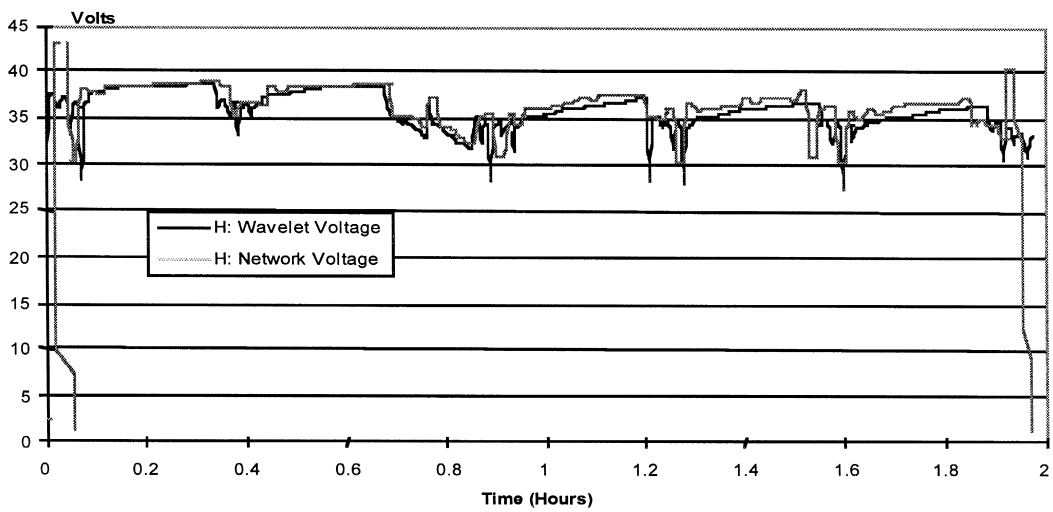

Fig. 12. Voltage response network predictions (gray) compared to the measured voltage response (black). The net was trained using test B (profile 1 at 20°C) with only 106 coefficients extracted using a wavelet representation of the original signal. The net was tested using the wavelet representation (106 coefficients) of test H (profile 2 at 0°C).

sponse of the cell. Again, these results reaffirm our conclusion advanced above, we believe that the net learned the correct temperature response of the cell.

10. Conclusions

Data representation one produced good results when trained with portions of both types of current profiles and portions of tests at each temperature. The amplitude/time representation (data representation one) required a little less than half the number of inputs (54 inputs) than the wavelet representation (data representation three, 106 inputs). However, the data pre-process is long. Representation one allowed the net to interpolate temperatures and patterns very well. But the net trained with representation one, when only one type of pattern and one temperature were used during the net training, had some problems extrapolating temperatures and pattern types (it underestimated them). The net showed good capabilities of extrapolating data over a period of 2 h when trained for just minutes with all the information available about temperature and pattern. We had only data measured over 2 h to compare with the net predictions. The net predictions obtained over 10 h agreed with our expectations based on experience and data available [3]. The net is capable of accurately predicting the output voltage with time for the lithium–ion battery.

Data representation two produced the correct voltage output for current patterns that it had not been trained with. However, the level of detail of the 1416 data point with 16 Fourier coefficients was not good enough to confirm that the net was capable of extrapolating behaviors (e.g., to predict temperatures outside the range of temperatures the net was trained with). The Fourier representation does not capture all the current spikes imposed on the cell, nor the voltage response spikes of the cell, but only a wave type response results. The Fourier representation did not require pre-processing of training data, as compared to data representation one. To better reproduce the original pattern, more Fourier coefficients would have to be used. The network should not experience any loss in performance in this case, but this increase in coefficients would mean an increase in input and output elements of the network. Another problem encountered with this representation was that the network could not be trained with part of a test and then tested with the rest of the same test, because the network had to be trained with Fourier coefficients from test segments of a similar length.

The wavelet representation (representation three) was similar to the Fourier representation in the sense that it required no pre-processing of the data before presentation to the network, unlike the amplitude/time representation (representation one). The wavelet coefficient data representation experienced the same problem as the Fourier representation, since the wavelet coefficients also had to be generated from test segments of the same length (in this research we used 2 h), i.e., the net cannot be trained in one segment of a current/voltage signature and tested in another segment of the same signature. However, the wavelet representation is more closely tuned to the current and voltage signatures. The spike behavior was easily captured; accordingly, we could assess the performance of the net trained with the wavelet representation than the net. Data representation three accurately reproduced the original voltage signal with a relatively small number of wavelet coefficients. Conversely, the discrete Fourier transform with 128 coefficients (including real and imaginary) did not reproduce more than small waves in response to steep voltage drops (spikes).

The net trained with the wavelet representation using only a given pattern and temperature predicted the response of the cell at different temperatures and current patterns imposed on the cells. The wavelet representation had the best extrapolation capabilities (and of course interpolation) when compared to the other representations. The wavelets data representation (data representation three) looks very promising in predicting the cycling failure of batteries, because it appears to capture the hidden functions describing temperature dependency and effect of long current plateaus (instead of spikes).

Coming back to the discussion of how many coefficients to choose to represent a given signature when Fourier or wavelet coefficients are chosen to represent the input current pattern, we can anticipate that if all of the sampling points were used, and more coefficients generated, greater detail in the signature would be produced, but at a cost of increasing the number of elements in the neural network. This increase in the number of elements is not a great concern, however, because although there may be many elements in an input vector, there will be only one input vector for each test used for training.

This result appears to indicate that the response to a current pattern is a cell property. A well-trained net will capture that mathematical function that correctly describes the cell's response to current patterns and temperature changes. That net will be able to extrapolate correctly the cell behaviour.

This paper does not represent an exhaustive study of the parameters that can affect the cycling life of a battery. We were limited to using only those parameters that were monitored during the tests used. If much information on variables was made available, a similar methodology to the one described in this paper could be used. For example, none of these data representations had any input elements representing the capacity loss due to repeated recharge/discharge cycles. This capacity loss is difficult to measure for such short tests because it is expected to be very small. Since the tests were conducted at different temperatures it is difficult to decide whether to attribute capacity loss to temperature change or repeated cycling. More extended databases are needed.

In each one of the cases presented, the network voltage predictions followed very closely the measured voltage pattern (when data were available to compare), and we can conclude that the predictive capabilities of the net trained with any of the three data representations are, as discussed, very good, making neural networks a very promising technique in predicting cycling life of batteries.

Acknowledgements

The authors are grateful for the support obtained from DOE/BES (Contract No. DE-FG02-93ER14387/A003). Dr. Urquidi-Macdonald is especially grateful to the NASA-ASEE summer fellowship program. Special thanks are extended to Dr. Chris Smith, SAFT America, for providing the data employed in this work (obtained under NASA contract).

References

[1] M. Urquidi-Macdonald, P. Egan, Author of a chapter in Corrosion Reviews, special issue on Applications of Computers in Corrosion, Vol. XV, Nos. 1–2. Freud Publishing House, 1997.
[2] Y.-H. Pao, Adaptive Pattern Recognition and Neural Networks. Addison-Wesley, 1989.
[3] A.L. Romero, Discharge Profile Test Report for the Pistol Grip Tool Engineering Battery, SAFT America, Cockeysville, MD. Internal Report Prepared for Fairchild Space, Washington Technical Support Center.
[4] G. Strand, T. Nguyen, Wavelets and Filters. Wellesley Cambridge Press, 1996.