Joule Counting Correction for Electric Vehicles Using Artificial Neural Networks

Michael D. Taylor

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue, 15213
Email: mdtaylor@cmu.edu

Abstract

Estimating the remaining energy in high-capacity electric vehicle batteries is essential to safe and efficient operation. Accurate estimation remains a major challenge, however, because battery state cannot be observed directly. In this paper, we demonstrate a method for estimating battery remaining energy using real data collected from the Charge Car electric vehicle. This new method relies on energy integration as an initial estimation step, which is then corrected using a neural net that learns how error accumulates from recent charge/discharge cycles. In this way, the algorithm is able to adapt to nonlinearities and variations that are difficult to model or characterize. On the collected dataset, this method is demonstrated to be accurate to within 2.5% to 5% of battery remaining energy, which equates to approximately 1 to 2 miles of residual range for the Charge Car given its 10kWh battery pack.

Introduction

Figure 1: The Charge Car electric vehicle

Importance of Battery Monitoring

One of the very difficult challenges involved in designing electric vehicles (EVs) is the accurate estimation of remaining energy in a partially-charged battery. Battery remaining energy (BRE) is important to quantify, as it is directly related to how far a vehicle may travel before requiring another charge. Because of the current size and cost of high-capacity lithium-based batteries, EVs typically have a maximum driving range between charges that is significantly less

than that of a car with a traditional internal combustion engine (ICE). For this reason, and because batteries cannot be instantly recharged or refueled like an ICE, quantifying BRE for range estimation is essential for trip planning.

Battery estimation in current generation electric vehicles is a major source of complaints from EV owners. Additionally, battery energy may be used to level, cool or heat the batteries, unlike in an ICE. This can cause confusion and distrust in consumers that do not understand these functions or cannot predict how this energy use (frequently referred to as "vampire drain") impacts their driving range.

Why Battery Monitoring is Difficult

Unlike the liquid fossil fuel in an ICE, whose energy content is directly proportional to its volume, the amount of energy contained within a battery cannot be measured directly, and must be inferred.

The amount of energy that can be drawn from a battery is determined by complex nonlinear relationships(Lee et al. 2008). Factors influencing the behavior of a battery include but are not limited to the state of charge, temperature of the cells and the rate of discharge. The temperature is determined by the battery's ambient environment and internal heat losses. The rate of discharge is typically measured in amperes, and also affects the terminal voltage of the battery. The rate of discharge is determined primarily by driving patterns, with losses increasing with more aggressive driving, uphill terrain, and high vehicle loads. This is because heat generation is primarily resistive in nature, and thus is increased by higher current loads. Temperature, current, and terminal voltage are measured by any typical battery monitoring system, as for the most part any effect of the world on the battery can be expressed by influencing these quantities.

Finally, because batteries are one of the more expensive components in an EV and not frequently replaced, they age over a period of several years, gradually losing capacity and the ability to discharge higher amounts of current for substantial amounts of time. This means that the relationships between temperature, voltage, discharge rate and BRE wll change over time, and furthermore, battery discharge behavior may vary from cell to cell, even among cells produced in the same production run.

Related Work

Several approaches for BRE and residual range estimation have been proposed. In general, while many approaches demonstrate high accuracy in simulation or using test-bench hardware with completely controlled parameters, relatively few have been shown to work on an actual electric vehicle. Furthermore, few approaches demonstrate the ability to adapt to changing vehicle parameters, and several are hard-coded with optimal parameters for one specific battery pack or cell.

Research into battery state estimation falls into three major categories: explicit physical modeling of battery behavior, adaptive methods that assume very little about the battery, and mixed methods which combine measured parameters with adaptive techniques such as extended Kalman filtering. Explicit battery modeling is beneficial in that it allows for accurate estimation of a specific battery pack or cell based on an understanding of real physical and chemical processes. Inexpensive battery testers may simply measure terminal voltage, labeling a battery as depleted once the terminal voltage drops below a certain voltage. The terminal voltage is only a partial indicator of state of charge, however, and changes depending on load current and temperature. For these reasons, open-circuit voltage is only used as a component in more complex models (Lee et al. 2008). Many of the constants required for modeling cannot be measured outside of very controlled conditions, however, and thus these methods usually lack the ability to adapt to different batteries or changes in battery condition. Adaptive methods, usually neural nets, exhibit the ability to approximate the nonlinear behavior of the battery and also have the potential to adapt as new data is collected during use, albeit without the one-to-one correspondence to physical properties of the battery. Mixed methods which implement simplified models along with extended or dual extended Kalman filters for correction exhibit some benefits of both methods, in that the models are grounded in reality but they are also able to adapt based on sensor uncertainty.

Ceraolo and Pede(Ceraolo and Pede 2001) demonstrate a model-based estimator for residual energy as well as residual range for electric vehicles. The battery model is based on an equation that describes the capacity of a lead-acid battery as a function of a constant discharge current and electrolyte temperature. The model requires known values for electrolyte characteristics such as freezing temperature as well as constants that must be empirically determined for the specific battery being considered. The battery is compared to a "fuel tank whose volume varies with the outlet flow rate." This model was implemented on an electric vehicle and achieved an accuracy of approximately +/- 5km range error over a distance of 80km. While demonstration on an actual vehicle is significant proof of concept, as the authors note, future work in the area must include the ability for the algorithm to adapt to variations in batteries, driver behavior, and weather.

Kim, Lee, and Cho(Kim, Lee, and Cho 2011) propose an improvement to the dual extended Kalman filter (DEKF) for SOC/capacity estimation. Initial model parameters for the DEKF are experimentally determined by discharging and charging the battery repeatedly at different temperatures. These charge/discharge profiles are matched with the battery behavior of the car at any point using a Hamming network, and the closest match is used to re-parameterize the DEKF. Checking the Hamming-predicted temperature against sensor data serves as a correction step. Benchtop charging and discharging tests yield an accuracy of +/-5% for SOC and capacity, but this publication does not yet implement the approach on an actual vehicle.

Ng et al.(Ng et al. 2009) demonstrate improvements to the coulomb counting method for SOC and SOH (state of health) estimation in small lithium ion batteries. The charging and discharging currents are integrated starting from an estimated SOC and assumed 100% SOH. Whenever the battery is fully discharged, the SOH is re-evaluated using the measured discharged capacity in Ah against the manufacturer-rated capacity. After a full discharge, once the battery is completely charged again, the efficiency coefficients of charging and discharging can be re-calculated based on the accumulated error in the system between full charges. When the batteries are discharged at a fixed current in a test setup, Ng et al. state an error rate of less than 1%. While this approach may be suitable for small batteries in consumer devices like cell phones, it is not particularly well suited to EVs because under normal operation, the car never fully drains its battery. Additionally, the load current varies dramatically from second to second (due to the driver's use of the accelerator, stop and go traffic, hills, etc.) and over larger time periods (due to changing drivers, changing weather conditions, etc.). Improved coulomb counting does, however, provide a reasonable estimate for energy discharged, and may be used as a feature or plant model as part of a more complex algorithm.

Qingsheng et al(Qingsheng and Chenghui 2010) simulate the use of an Elman neural network for SOC estimation. (A similar method is discussed in (Rui-hao, Yu-kun, and Xiao-fu 2011).) The training and testing data are taken from the ADVISOR database of drive data. The Elman NN structure includes a context layer that uses the outputs of the hidden layer at the last time step as additional inputs in the present hidden layer. This allows the network to retain a single instance of historical data, which allows it to identify increasing or decreasing inputs. The average percent error for SOC is stated to be approximately 1.2%, though according to the figures the test data appears to only cover the top 2% of the full battery capacity. It is unknown how these method perform for batteries that are moderately or severely drained, either in a benchtop setting or on an actual EV.

Shen et al.(Shen and Chau 2005) propose a unique variation on the neural net estimation approach where each of four inputs is assigned a value for the total discharged current (Ah) within fixed non-overlapping ranges. Two additional inputs are assigned to the regenerated capacity and the pack temperature. The rationale for segmenting expended capacity by the magnitude of the current when it was drained comes from the fact that a battery drained at high current has a lesser capacity than if it were drained at low current. This approach, when trained and tested in a benchtop setup with high accuracy data acquisition tools, yields an average per-

cent error of 2.67%. Again, it would be very informative to see how this approach performs using actual EV data.

Research Goals

In this paper, we seek to develop a method for estimating battery remaining energy that achieves the following goals by combining the best features of previously proposed methods.

Explicit Error Modeling Given the significant effect of energy loss as heat for EV batteries, our method should model the energy lost in order to present the user with a more accurate figure for BRE.

Responsive Training Given that EV battery packs each have unique charge and discharge behaviors, the method should be capable of being quickly trained online using real EV drive data.

Adaptation Given that drivers may change, temperature rises and falls with the seasons, and the car and battery packs may age and degrade, the method should constantly re-train and evolve to deliver acceptable estimates for the current set of conditions.

Data Summary

In the following experiments, we use a dataset collected from the Charge Car electric vehicle (shown in fig. 1). The vehicle is equipped with a 10kWh bank of lithium iron phosphate batteries, with an estimated range of 40 miles per full charge. The dataset was collected by four independent drivers over the course of 81 days, logging a total of 51 hours of drive time over 710 miles. This accompanies an expenditure of 177 kWh over approximately 15 full charge cycles. The drivers were instructed to drive normally without charging until the battery was as close to depletion as possible. In this way, we ensured that we have collected representative data samples across the full range of battery states. Given the lack of effective implemented battery estimation prior to the development of this method, the drivers relied on pack voltage alone, so the actual range driven during each cycle varied. As expected, more data exists for nearly-charged states than nearly-depleted states, as shown in fig. 2.
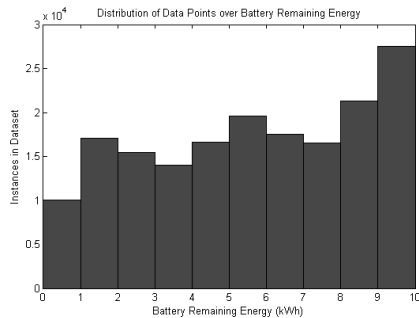
Figure 2: Distribution of data samples by discharge level. Note that we have data at both fully charged and nearly discharged states.

Linear Fit Method

Prior research (add citations) suggests that battery behavior is highly non-linear. We justify this on our own electric vehicle by attempting to create a linear combination of temperature, pack voltage, charge and load current, and charge and load power. For this model, we derive our constants for the linear fit using least-squares methods on historical data. Fig. 3 demonstrates the performance of the linear fit across the entire dataset, where the parameters are discovered using least-squares over the first half of the data. Because the instantaneous sensor values for current, temperature, and voltage change drastically from second to second, the fit is rather poor, again suggesting that a method for either encoding non-instantaneous (smoothed) features or nonlinearities is needed.
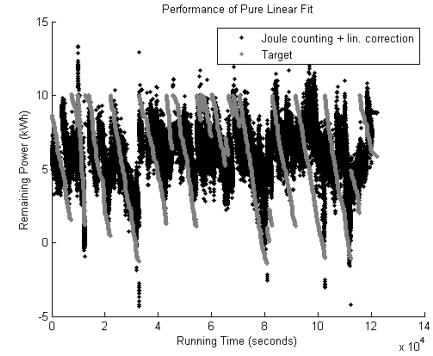
Figure 3: Linear fit of sensor values to remaining energy, where parameters are fit to the first half of the dataset. The performance is erratic and noisy.

Correction Factor Algorithm

We identify three general states for an electric vehicle, as shown in fig. 4. In the driving and charging states, the vehicle collects time-stamped battery sensor data while providing an online estimate of remaining energy. In the fully charged state, the vehicle corrects for accumulated error and adapts to changes in vehicle or driver behavior. Combining techniques from coulomb counting and neural nets yields an estimation method that quickly adapts to a given electric vehicle, driver, and environment in order to provide a reliable estimate for the amount of energy remaining in the EV battery. The algorithm relies on joule counting (which tracks energy in kWh as opposed to charge in coulombs) from the most recent full charge for its initial estimate. The neural net component then produces an estimate for the error accumulated by joule counting, which is used to adjust the instantaneous BRE estimate. The neural net output is low-pass filtered to produce a smooth, continuous estimate to the driver. This process begins again after every full charge cycle.

Joule Counting Estimate

In order to calculate an initial BRE estimate, we use a modified joule counting approach where instead of simply integrating current over time (Ah), we integrate power over time
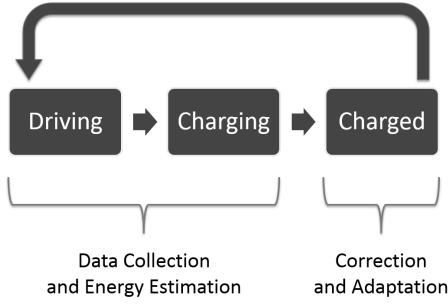
Figure 4: Depiction of one drive cycle, which includes a charging and discharging phase.

(kWh). This is because we expect residual energy (rather than residual charge) to be more useful in estimating the remaining available range in an EV, given the large variations in both voltage and current due to the electric motor load. In this sense, it is more appropriate to refer to this as Joule counting.

We must first identify states where we assume the battery to be fully charged. We do this by looking for data points in the un-cut dataset flagged as charging or parked where either the source current is at a low trickle-charge or the last instance of non-zero current is at trickle-charge levels (suggesting that the car has been unplugged after a full charge but not driven). At each of these points, we reset the BRE to a fixed quantity equal to the nominal battery capacity of the pack. (We assume that the battery capacity does not degrade over the course of our dataset, and we will later suggest future directions for estimating a non-static battery capacity.) At every second, with the arrival of a new data point, we integrate the power leaving and entering the battery as the product of pack voltage and current with respect to time, according to equation 1. This estimate is only a rough approximation of BRE, as it does not account for internal losses and will have accumulated a non-zero error by the next instance of full charge.

$$E_R(t_{now}) = E_{BC} - \int_{t_{charged}}^{t_{now}} I(t)V(t)dt \qquad (1)$$

Joule Counting Correction

Nominally, if a car expends a number of kilowatt hours while driving, it should be expected that it should take the same number of kilowatt hours for it to reach a fully charged state. In practice, our data shows that the integral of power between two fully charged states is non-zero. This quantity can be calculated by equation 2. This result is comparable to the results demonstrated in (Ng et al. 2009). Because we can detect the fully-charged states, however, and because we know the approximate shape of the BRE between these states, we can distribute the accumulated error post-facto in order to establish a target for the neural net correction factor. As shown in fig. 6, the error accumulates at an average rate of 100 watts over the course of the dataset if left uncorrected.

$$E_{err} = \int_{t_{charged}}^{t_{now}} P_{err}(t)dt = \int_{t_{charged}}^{t_{now}} I(t)^2Rdt \qquad (2)$$

The energy error across one discharge-charge cycle can be expected to accumulate gradually, and we must choose a way to distribute it. The simplest means is to redistribute the error between charges linearly with time. Two more intuitive and potentially more realistic methods include redistributing error according to either current or current squared (I^2), which suggests that error accumulates the most when the car is in heavy use. The goal of redistributing the error is to continuously correct the historical residual energy estimate so that the net energy spent and regained is zero between charges. The errors are small, so except for occasional outlier drives, we find that different correction methods yield very similar results. The histogram in fig. 5 also confirms that the difference between linear and I^2 correction is centered tightly on 0. In this paper, we choose I^2 as our historical correction method because it matches a physical model of resistive heating, which is proportional to current squared. Resistive heating in the batteries cannot be measured, and thus is one of the major sources of energy loss error.
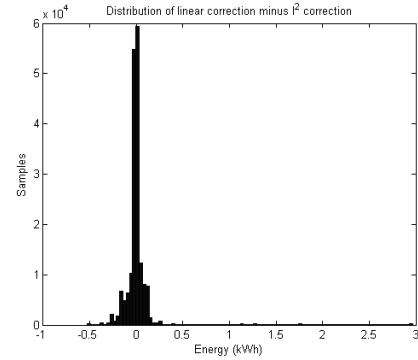

Figure 5: Histogram showing that nearly all of the data points for linear time correction and I^2 correction deviate by less than .25 kWh.

By calculating the energy error and the integral of current squared between the two charge cycles, we can calculate the fixed virtual resistance R for that interval that corrects for the error (equation 3), and then calculate the exact error correction in kilowatts for each instance between the charge cycles using equation 4. In this way, the corrected estimate between two fully charged instances will contain no discontinuities, as shown in fig. 7

$$R = \frac{E_{err}}{\int_{t_{charged}}^{t_{now}} I(t)^2dt} \qquad (3)$$

$$P_{GT}(t) = P_{JC}(t) + P_{err}(t) = P_{JC}(t) + I(t)^2R_{err} \qquad (4)$$
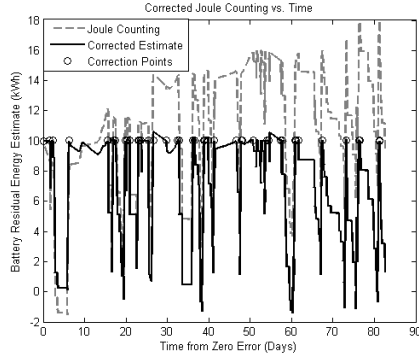
Figure 6: Plot showing joule counting across the entire dataset. Notice the upward error trend in the uncorrected estimate.
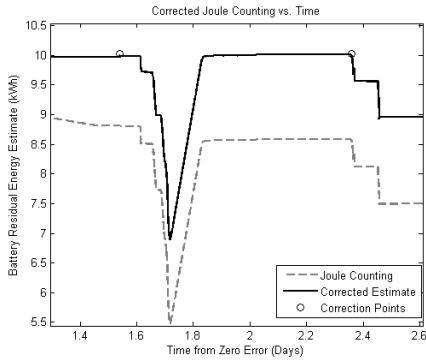

Figure 7: Section of joule counting estimate showing two correction points. Note the increasing error in the uncorrected estimate.

The Need for a Learning Correction Factor

One proposed method for estimating battery remaining energy is to use the correction constant from the previous historical correction in real-time on subsequent drives. For the I^2 correction method, this is analogous to estimating the internal resistance of the battery after each drive cycle. As fig 8 shows, however, estimating internal resistance based on the median of past values cannot be used for prediction. Using only the most recent value results in even more erratic performance. Thus, we hypothesize that the correction constants are not accurately estimating the internal resistance of the battery. The variation suggests that the correction factor is not constant, and encodes non-linear behavior as a factor of temperature and time-variant properties. Because of this, a learning-based approach capable of capturing these non-linearities may be better suited to real-time correction.

Neural Net Correction Factor

The error correction previously described can only be calculated for instances in the past, between two fully-charged instances that have already occurred. Therefore, we must estimate at any given instant what we expect the joule counting
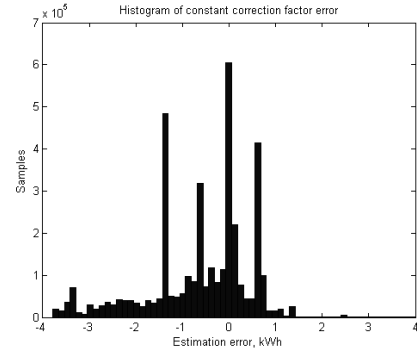

Figure 8: Error from I^2 ground truth using prediction based on prio correction factor values, with errors of several kW.

error to be. For this purpose, we use a neural net with one hidden layer of 25 neurons. The neural net has inputs corresponding to the instantaneous pack voltage, current, and average cell temperature, as well as averages over the past several seconds. It also utilizes the raw joule-counted estimate (reset at the last full charge) as input.

After each full charge, the neural net re-trains itself across an interval of the most recent historical data equal to the length of three full charge/discharge cycles, or seven hours of drive time. Batch training is used with 70% training points, 15% testing points, and 15% validation points, selected over the past interval. The weights for the nets are initialized to their previous values, which allows for some historical influence on the net while simultaneously allowing it to adapt primarily to the most recent drives.

During operation, the output of the neural net given instantaneous sensor values is added to the raw joule-counted estimate to yield a corrected instantaneous estimate for BRE, shown in fig. 9. If we remove the chronology of the estimate and simply compare the NN-corrected joule counting estimate against the post-facto ground truth estimate, as in fig. 10, we see that our estimate fits very closely.

$$E_{JC}(t_{now}) = E_{BC} - \int_{t_{charged}}^{t_{now}} I(t)V(t)dt \qquad (5)$$

A histogram of errors when compared to ground truth (fig. 11) reveals that for this particular training instance, the neural net has a tendency to overestimate the correction factor. The error shown here is from -.25 to +.5 kWh, which corresponds to approximately -1 to +2 miles of range error for the Charge Car EV based on a maximum range of 40 miles. Fig. 12 shows that the error as well as the variance for our estimation decreases with additional training events. Because the error is still decreasing after all of the data has been trained and processed, additional data should be used to verify the asymptotic error limit. The author expects, however, that the error cannot substantially decrease beyond the range shown.

Alternative Linear Correction Factor

After characterizing the neural net correction factor, we revisited the notion of linear fits by formulating a similar cor-
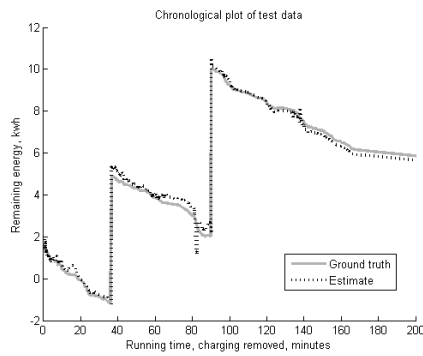
Figure 9: Chronological plot of ground truth versus corrected joule counting, showing DC agreement.
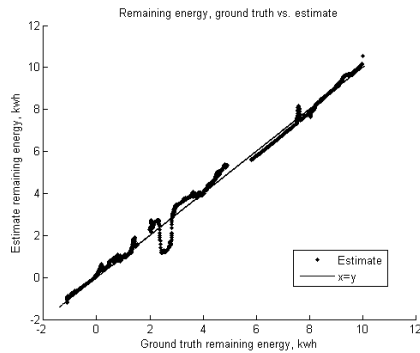

Figure 10: Plot of ground truth estimates versus corrected joule counting estimates. Note the general agreement, and outliers which can be corrected in smoothing.
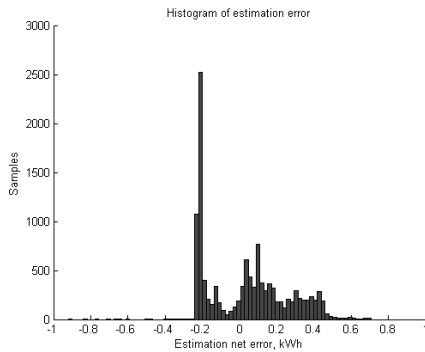

Figure 11: Histogram of algorithm error in kWh, displaying a slight positive bias.

rection factor using linear fits rather than the neural net. Here, we employ a similar least-squares fitting method, but instead of fitting to the historical corrected remaining energy, we fit to the accumulated historical error over the first half of the dataset as corrected using the I^2 method. In fig. 13, we see that the linear correction factor adequately predicts the remaining energy with similar performance to the
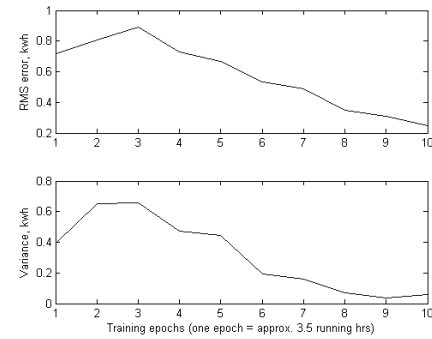

Figure 12: Error and variance of the algorithm by training epoch, generally decreasing with additional training data.
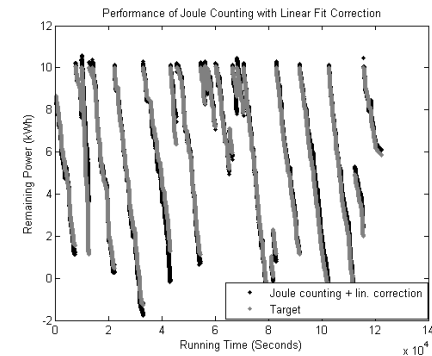
neural net correction factor.

Figure 13: Using a linear fit to the accumulated error, fit to the first half of the dataset, yields a reasonable estimate for remaining energy.

Comparison of Correction Factors

On closer inspection in fig. 14, we see that the neural net correction factor matches X=Y better, but presents outliers that deviate more than the linear correction factor. This suggests that the linear correction factor might be more appropriate in scenarios where outliers such as those shown are less tolerable.

Output Smoothing

One significant complaint among users of electric vehicles is that the range estimates given by their cars tend to vary from minute to minute. Given that the output of the neural net depends largely on instantaneous sensor readings and does not carry any information from its previous state, its output is highly stochastic due to sensor noise and neural net complexity. As such, outliers have the potential to mislead drivers. One simple method for smoothing the output estimate would be to average the full estimate over the past several minutes. This introduces a time delay, however, and thus interferes with the drivers' interpretation of the estimate. Instead, we average only the correction factor from
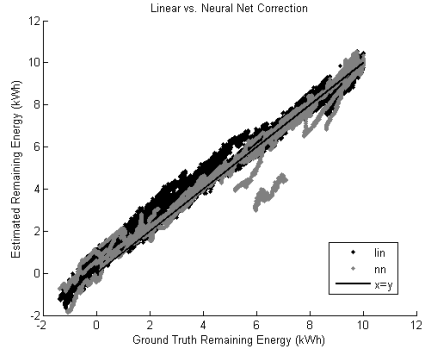
Figure 14: Comparison of linear fit and neural net correction factors. Neural net correction factors are more tightly grouped around x=y, but also present some outliers not present in the linear fit correction factors.

the neural net. We expect the error to accumulate at a significantly lower rate than energy is expended, so the time delay introduced by averaging the correction factor is negligible. This correction is expressed in equation 6.

$$E_{est}(t) = E_{JC}(t) + \frac{1}{n+1} \sum_{t=t_{now}}^{t_{now}-n} E_{NN}(n) \qquad (6)$$

As can be seen in figs. 15 and 16, the smoothed estimate eliminates some of the discontinuities and more egregious errors yielded by the instantaneous estimate. This is particularly notable in fig. 16 at approximately 2.5 kWh BRE.
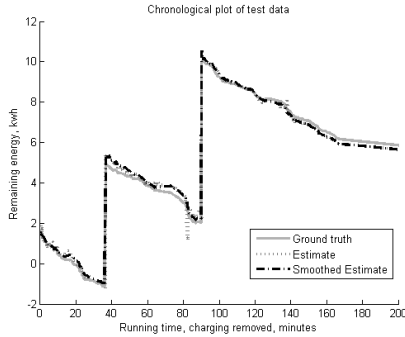

Figure 15: Chronological plot of smoothed estimate and ground truth. Note that the outliers present in the un-smoothed estimate have been significantly reduced.

Conclusions and Future Directions

In this paper, we have demonstrated a new method for estimating battery remaining energy based on energy integration with a corrective neural net that allows for rapid adaptation to battery nonlinearities as well as changes in drivers and driving conditions. Furthermore, the method is trained and tested using real data collected from the Charge Car electric vehicle, thus strongly supporting its feasibility in actual EVs. The RMS error for this method is approximately
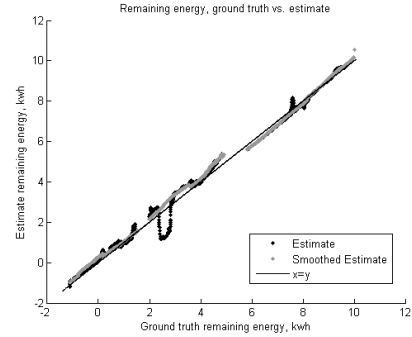

Figure 16: Smoothed estimate versus ground truth, showing significant outlier reduction.

2.5% to 5% of the total battery capacity after significant training, which translates to approximately 1 to 2 miles in range given the 10 kWh battery pack in the Charge Car.

After the method was implemented, there is some concern about overfitting of the neural net, since the training and testing data are randomly selected from the training interval, rather than selected from separate drives. This likely does lead to overfitting to the most recent drives, though given the acceptable performance of the real-time estimate, we do not believe that this overfitting causes serious negative consequences to the methods performance. In subsequent iterations, selecting the training interval to exactly coincide with recent drives instead of drive time, and selecting training and testing data from different drives may yield better performance. It may also yield worse performance if the one of the drives selected for training or testing is a significant outlier (caused, for example, by a new driver taking a single short drive but not interacting with the vehicle any further).

Future testing of this method could be performed on a larger dataset to better characterize its adaptation to changing seasons, new drivers, and battery age. Additionally, it could be implemented on a live vehicle using a small in-vehicle PC, which would ideally prove its capability for online learning. To further the usefulness of the estimator, maximum battery capacity should also be estimated as the battery ages over time, and residual range could additionally be predicted using additional features from the electric vehicle, such as transmission state, GPS data, and wheel RPM.

Battery estimation is relevant not only to electric vehicles but also to mobile devices, backup power supplies, and other devices where electrical energy reserves are used. While the method for estimation may vary between these domains (Meissner and Richter 2003), they may provide additional data that could be used for testing the joule counting correction method, and where processing power is available, this method may prove beneficial over the state of the art.

Car team as well as the Charge Car sponsors who made this work possible by providing the data and test platform for this research.

References

Ceraolo, M., and Pede, G. 2001. Techniques for estimating the residual range of an electric vehicle. *IEEE Transactions on Vehicular Technology* 50(1):109–115.

Kim, J.; Lee, S.; and Cho, B. 2011. SOC/capacity estimation of a Li-Ion battery based on pattern recognition at different temperature. *8th International Conference on Power Electronics - ECCE Asia* 296–303.

Lee, S.; Kim, J.; Lee, J.; and Cho, B. 2008. State-of-charge and capacity estimation of lithium-ion battery using a new open-circuit voltage versus state-of-charge. *Journal of Power Sources* 185(2):1367–1373.

Meissner, E., and Richter, G. 2003. Battery Monitoring and Electrical Energy Management. *Journal of Power Sources* 116(1-2):79–98.

Ng, K. S.; Moo, C.-S.; Chen, Y.-P.; and Hsieh, Y.-C. 2009. Enhanced coulomb counting method for estimating state-of-charge and state-of-health of lithium-ion batteries. *Applied Energy* 86(9):1506–1511.

Qingsheng, S., and Chenghui, Z. 2010. Battery State-Of-Charge estimation in Electric Vehicle using Elman neural network method. *... (CCC), 2010 29th ...* 5999–6003.

Rui-hao, L.; Yu-kun, S.; and Xiao-fu, J. 2011. Battery state of charge estimation for electric vehicle based on neural network. *2011 IEEE 3rd International Conference on Communication Software and Networks* 493–496.

Shen, W., and Chau, K. 2005. Neural network-based residual capacity indicator for nickel-metal hydride batteries in electric vehicles. *..., IEEE Transactions on* 54(5):1705–1712.