

Learning Policies For Battery Usage Optimization in Electric Vehicles

Stefano Ermon, Yexiang Xue, Carla Gomes, and Bart Selman

Department of Computer Science
Cornell University
Ithaca, NY

{ermonste,yexiang,gomes,selman}@cs.cornell.edu

Abstract. The high cost, limited capacity, and long recharge time of batteries pose a number of obstacles for the widespread adoption of electric vehicles. Multi-battery systems that combine a standard battery with supercapacitors are currently one of the most promising ways to increase battery lifespan and reduce operating costs. However, their performance crucially depends on how they are designed and operated.

In this paper, we formalize the problem of optimizing real-time energy management of multi-battery systems as a stochastic planning problem, and we propose a novel solution based on a combination of optimization, machine learning and data-mining techniques. We evaluate the performance of our intelligent energy management system on various large datasets of commuter trips crowdsourced in the United States. We show that our policy significantly outperforms the leading algorithms that were previously proposed as part of an open algorithmic challenge.

1 Introduction

Electric vehicles, partially or fully powered by batteries, are one of the most promising directions towards a more sustainable transportation system. However, the high costs, limited capacities, and long recharge times of batteries pose a number of obstacles for their widespread adoption. Several researchers in the field of Computational Sustainability [1] have addressed aspects of this problem. In particular, there is an active line of research focusing on improving navigation systems with novel routing algorithms, both by taking into account specific features of electric vehicles [2], and by considering new aspects such as real-time information about road conditions and traffic lights [3].

In this paper, we focus on a complementary aspect of the problem that is optimizing the energy efficiency of batteries in electric vehicles. There are two main sources of inefficiencies in batteries. The first one is that due to internal resistance, battery energy is partially wasted as heat when it is charged and discharged. The second one is that due to Peukert's Law, the actual delivered capacity of a battery depends on the rate at which it is discharged. Furthermore, current battery technology imposes rather severe limits on the number of charge/recharge cycles a battery can handle, thus reducing their lifespan and increasing operating costs.

One promising direction towards addressing these issues are multi-battery systems, such as the ones proposed in [4] and [5], which integrate a standard battery with one or more supercapacitors, as depicted in Figure 1. Intuitively, the idea is that while the battery is good at holding the charge for long times, the supercapacitor is efficient for rapid cycles of charge and discharge. Using the capacitor as an energy buffer, one can significantly increase the battery’s lifespan by reducing its duty. In fact, although supercapacitors have low energy densities, they behave like an ideal battery that can efficiently handle millions of full charge/discharge cycles. The performance of these systems is heavily dependent on how they are managed. In this direction, there has been recent work in the automated planning community on the optimal scheduling of multi-battery systems [6, 7]. However, previous work assumes full knowledge of an underlying probabilistic model describing the system, which is not available for electric vehicles. Since it would be very difficult to construct such a model using a priori information, we take a data driven approach to the problem. In particular, we leverage a large dataset of commuter trips collected across the United States by Chargecar [8], a crowdsourcing project open to the public, and we construct an efficient management scheme using a sample-based optimization approach. Specifically, after defining a suitable set of features, we learn an empirical Markov Decision Process (MDP) model from the available data, and we compute a policy that optimizes the average performance. This policy is represented as a large table of state-action pairs, and is only defined for states that were previously observed in the dataset, while we wish to construct a more general management scheme that applies to a wider range of scenarios. We therefore use this policy as a training set, and we use supervised learning techniques to learn a new policy that compactly represents the information available and generalizes to situations previously unseen in the training set. This policy is shown to outperform the leading algorithms that were previously proposed as part of an open algorithmic challenge.

2 Sampling-based Optimization

We consider a probabilistic planning problem formulated as a Markov Decision Process (MDP). An MDP is a tuple (S, A, P, c) where S is a set of states, A is a set of actions, P is a set of transition probabilities and $c : S \times A \times S \mapsto \mathbb{R}$ is an (immediate) cost function. If an agent executes an action $a \in A$ while in a state $s \in S$, then it incurs in an immediate cost $c(s, a, s')$ and it transitions to a new state $s' \in S$ with probability $P(s'|s, a)$. We denote by $\mathcal{A}(s) \subseteq A$ the set of actions available while the agent is in state s . Further, there exists a finite set of *goal states* $G \subseteq S$, where the agent stops to execute actions, and no longer incurs in any cost.

In this paper, we consider a class of factored MDPs where $S = X \times Y$ so that any state $s \in S$ has two components $s = (x, y)$ with $x \in X$ and $y \in Y$. We assume the dynamics of the two components are independent, i.e. the transition probabilities can be factored as follows

$$P((x', y')|(x, y), a) = P_x(x'|x)P_y(y'|y, a).$$

Notice that x' does not depend on the action a , which only affects the component y' . The two components of the state are however coupled by the immediate cost function

$c((x, y), a, (x', y'))$, which can depend on both x and y . We assume that $P_y(\cdot)$ and the immediate cost function $c(\cdot)$ are known but $P_x(\cdot)$ is unknown. However, we are given a set of K i.i.d. sample trajectories $\mathcal{T}_1, \dots, \mathcal{T}_K$ of the x component of the state space, where

$$\mathcal{T}_i = (x_0^i, x_1^i, \dots, x_{T_i-1}^i)$$

is sampled according to $P_x(x'|x)$ and $x_{T_i-1}^i \in G_x$ is a goal state. For example, in our battery management application, each trajectory corresponds to one commuter trip. Given this information, our objective is to find an admissible policy that minimizes the expected cost for this partially unknown MDP.

Given $x, x' \in X$, let $f(x, x')$ be the empirical transition probability from x to x' according to the available samples (the number of times x' appears immediately after x over the number of times x appears). We can define DP equations based on the sampled transition probabilities as follows

$$V(x, y) = \min_{a \in \mathcal{A}(x, y)} \left(\sum_{x' \in X} \sum_{y' \in Y} f(x, x') P_y(y'|y, a) (c((x, y), a, (x', y')) + V(x', y')) \right)$$

for all observed states $x \in \bigcup \mathcal{T}_i$. Solving the DP equations, we can compute the “optimal posterior action” $a^*(s) = a^*(x, y)$ for all $x \in \mathcal{T}_i$ and for all $y \in Y$, that is the action minimizing the total expected cost according to our maximum-likelihood estimate of the underlying MDP model. Notice that the “optimal posterior action” $a^*(s)$ converges to the true optimal action for the MDP as $K \rightarrow \infty$ because $f(x, x') \rightarrow P_x(x'|x)$ (assuming the initial states x_0^i are uniformly sampled).

Although the number of distinct states $x \in \bigcup \mathcal{T}_i \subseteq X$ can be very large, the samples $\mathcal{T}_1, \dots, \mathcal{T}_K$ do not necessarily cover the entire state space X . We therefore wish to obtain a compact representation of the policy $a^*(\cdot)$, that hopefully will be able to generalize to states $x \in X$ such that $x \notin \bigcup \mathcal{T}_i$, i.e. states previously unseen in the set of available samples. We therefore generate a labeled training set of state-action pairs

$$\bigcup_i \{((x, y), a^*(s)), x \in \mathcal{T}_i, y \in Y\}$$

and we use supervised learning to learn a policy $\pi : S \rightarrow A$. Notice that the particular structure of the problem, with independent dynamics and partially known transition probabilities, allows us to artificially generate $|Y|$ time more training examples than what we originally started with. This aspect leads to significant improvements in our battery management application problem.

2.1 Related Work

Sampling-based solutions, where a finite number of sampled trajectories is used to optimize performance, are a popular technique in the fields of approximate dynamic programming [9] and reinforcement learning [10], especially for complicated systems for which a model is not known or only available through simulation. However, unlike Reinforcement Learning we are dealing with a non-interactive learning scenario, where

we cannot choose how to interact with the system while samples are collected. Specifically, the learning process occurs offline and in batch. Further, since we have partial knowledge about the MDP (e.g., the immediate cost function), we use a model-based method [11] similar to the *certainty equivalence method* [12] where we estimate the most likely transition probabilities for the unknown part of the model. Unlike the dominant approach that uses function approximations to represent the value function (or Q-values) [13–15] and selects the action based on the greedy policy with respect to the estimated values, we directly represent a policy mapping states to actions. Specifically, we use supervised learning to train a policy using a dataset of “posterior optimal” actions computed according to the learned MDP model. The policy compactly represents the available information and, in our application, empirically performs better than models fitted to the Q-values. Further, in this way we can directly analyze the structure of policy being used, thus simplifying the deployment on a vehicle.

Our approach is similar to a line of research in the planning community [6, 16–18], where researchers have tried to learn strategies to solve planning problems in a fixed domain by learning from particular solutions to training problems sampled from the same domain. Specifically, our work is most closely related to [6], where they use a sample-based approach to learn a policy for a multiple-battery system modeled as an MDP. However, since we are dealing with electric vehicles we need to optimize charge/discharge cycles while they focus on the discharge aspect only. Consequently, we use a quadratic objecting function, while they optimize for the plan length. Further, their work is based on synthetic data generated from a known model, while we face the problem of learning a probabilistic model from real-world crowdsourced data, which creates additional challenges such as feature selection. Furthermore, in [6] they use as training examples sequences of state-action pairs where the actions are obtained optimizing in hindsight for a single sample realization of the randomness, i.e. the training set is generated using an optimal omniscient policy that knows the future ahead of time. Although the method is shown to perform well in practice, it doesn’t provide any theoretical guarantee. As a counterexample, consider a simple MDP modeling a lottery with an expected negative return, where the actions are either bet on a number or not to play at all. Given any realization of the randomness, the optimal omniscient policy would always suggest to bet (since it knows the future), but the optimal risk-neutral strategy is not to play the lottery, and therefore it cannot be learned from such training examples. In contrast, we also use a form of hindsight optimization, but we jointly consider all the samples, using a sample-based approximation for the expectation that provably converges to the true optimal value as the number of samples grows.

In the remainder of the paper, we will describe how we apply this general approach to the battery management application.

3 Problem Description

There are two main sources of inefficiencies in batteries. The first one is that batteries have an internal resistance R_{int} , and therefore they dissipate power as heat as $R_{int}i^2$ when charged or discharged with a current i . Secondly, the capacity of a battery is related to the rate at which it is discharged by Peukert’s Law. In particular, the faster a




Fig. 1: Architecture of the battery system and sign convention used (a positive number indicates current flowing in the direction of the arrow).

battery is discharged with respect to the nominal rate (by pulling out a higher current), the smaller the actual delivered capacity is. The effect depends on the chemical properties of the battery and is exponential in the size of the current. Therefore, substantial savings can be obtained by reducing the current output from the battery used to achieve a certain desired power.

One promising direction towards improving battery efficiency are multiple-battery systems such as the ones proposed in [4] and [5], which integrate a standard battery with one or more supercapacitors, as depicted in Figure 1. Intuitively, the idea is that the battery is good at holding the charge for long times, while the supercapacitor is efficient for rapid cycles of charge and discharge. Using the supercapacitor as a buffer, high peaks in the battery's charge and discharge currents can be reduced, thus reducing the losses due to Peukert's Law and the internal resistance. In fact, supercapacitors are not affected by Peukert's Law and behave like an ideal battery. Furthermore, this can substantially increase the lifespan of batteries because of the reduced number of full charge-discharge cycles the battery must handle (supercapacitors on the other hand can handle millions of full charge/discharge cycles). Improvements in battery efficiency lead to reduced costs, increased range, and therefore more practical electric vehicles.

While the savings obtained with multi-battery systems can be substantial, they heavily depend on how the system is managed, i.e. on the strategy used to charge and discharge the capacitor. Managing such systems is non-trivial because there is a mix of vehicle acceleration and regenerative braking (when power can be stored in the battery system) over time, and there is a constraint on the maximal charge the capacitor can hold. For instance, keeping the capacitor close to full capacity would allow the system to be ready for sudden accelerations, but it might not be optimal because there might not be enough space left to hold regenerative braking energy. Intelligent management algorithms therefore need to analyze driving behavior and vehicle conditions (speed, acceleration, road conditions,...) in order to make informed decisions on how to allocate the power demand. Intuitively, the system needs to be able to predict future high-current events, preparing the supercapacitor to handle them and thus reducing the energy losses on the battery. The results can be quite impressive. In figure 2 we show the battery output over a 2 minutes window of a real world trip, when no capacitor is used, when it is managed with a naive buffer policy (charging the capacitor only from regenerative

braking, and utilizing the capacitor whenever there is demand and energy available), and when it is managed by our novel system DPDecTree. While the total power output is the same in all 3 cases, when the system is managed by DPDecTree the output tends to be more constant over time, thus reducing the energy wasted due to battery inefficiencies.




Fig. 2: Battery power output over time. A smoother more constant-like curve means reduced losses due to battery inefficiencies.

4 Modeling

To formalize the battery management problem described earlier, we consider a discrete time problem where decisions need to be taken every δ_t seconds. A decision is a choice for the variables (i_{bc}, i_{bm}, i_{cm}) in Figure 1, where i_{bc} is the current flowing from the battery to the capacitor, i_{bm} and i_{cm} are the currents from the battery and capacitor to the motor, respectively. These variables must satisfy certain constraints, namely the capacitor cannot be overcharged or overdrawn and the energy balance must be preserved. As a performance metric, we consider the i^2 -score proposed in [4] and used in the Chargecar contest [8], where the objective is to minimize the sum of the squared battery output current $(i_{bc} + i_{bm})^2$ over time. Intuitively, reducing the i^2 -score means reducing the energy wasted as heat and due to Peukert’s Law, as well as increasing battery lifespan [19].

We first consider a simplified setting where we assume to know the future energy demand of the motor (positive when energy is required for accelerations, negative when energy is recovered from regenerative braking) ahead of time. This translates into a deterministic planning problem because we assume there is no more randomness involved. By computing the optimal sequence of actions (since the problem is deterministic, we don’t need a policy), we obtain a lower bound on the i^2 -score that is achievable in real-world problems where the future demand is not known.

4.1 A Quadratic Programming Formulation

Consider a single trip, where T is the number of discrete time steps (of size δ_t) in the control horizon. Let C_{max} be the maximum charge the capacitor can hold. As previously noted in [20], the problem can be formalized as a Quadratic Program. Specifically, we wish to minimize

$$\min \sum_{t=0}^{T-1} (i_{bc}(t) + i_{bm}(t))^2$$

subject to

$$i_{cm}(t) + i_{bm}(t) = d(t), \forall t = 0, \dots, T-1 \quad (1)$$

$$0 \leq \sum_{k=0}^t i_{bc}(k) - i_{cm}(k) \leq C_{max}/\delta_t, \forall t = 0, \dots, T-1 \quad (2)$$

where $d(t)$ is the motor demand at time step t . The first set of constraints (1) requires that the demand $d(t)$ is met at every time step $t = 0, \dots, T-1$. The second set of constraints (2) ensures that the capacitor is never overcharged or overdrawn (the capacitor is assumed to be empty at the beginning of the trip, and not to lose charge over time). Notice that the battery charge level over time is completely determined by the decision variables, and does not affect the i^2 -score.

Reducing the dimensionality We introduce a new set of variables

$$\Delta(t) = i_{bc}(t) - i_{cm}(t), t = 0, \dots, T-1$$

and using (1) we can rewrite the objective function as

$$\sum_{t=0}^{T-1} (i_{bc}(t) + i_{bm}(t))^2 = \sum_{t=0}^{T-1} (i_{bc}(t) + d(t) - i_{cm}(t))^2 = \sum_{t=0}^{T-1} (\Delta(t) + d(t))^2$$

Further, the constraints (2) can be rewritten in terms of $\Delta(t)$ as

$$0 \leq \sum_{k=0}^t \Delta(k) \leq C_{max}/\delta_t, \forall t = 0, \dots, T-1 \quad (3)$$

In this way we have simplified the problem from $3T$ variables $\{(i_{bc}(t), i_{bm}(t), i_{cm}(t)), t = 0, \dots, T-1\}$ to T variables $\{\Delta(t), t = 0, \dots, T-1\}$.

The resulting Quadratic Programs can be solved to optimality using standard convex optimization packages, but long trips can take a significant amount of time (see comparison below). Since we will later consider the stochastic version of the planning problem, we consider an alternative approximate solution technique that takes into account the sequential nature of the problem and generalizes to the stochastic setting.

4.2 A Dynamic Programming Solution

A faster but approximate solution can be obtained via Dynamic Programming by discretizing the capacity levels of the supercapacitor with a step δ_C and then recursively solving the Bellman equation

$$J(t, C) = \min_{0 \leq C' < N} \left\{ (d(t) + \delta_C / \delta_t (C' - C))^2 + J(t + 1, C') \right\}$$

for $t = 0, \dots, T - 1$, with boundary condition

$$J(T, C) = 0 \forall C$$

If the maximum capacity C_{max} is discretized with N steps, the complexity is $O(N^2T)$ for a trip of length T . This method is faster than solving the previous Quadratic Program directly, even though the solution is suboptimal because of the discretization. Furthermore, the DP solution does not just provide a sequence of “optimal” actions, but an actual policy that gives the action as a function of the current capacity level (for a fixed load profile).

Choosing the discretization step There is a tradeoff involved in the choice of the discretization step δ_C of the capacity level. The smaller δ_C is, the better is our approximation to the original QP, but the running time also grows quadratically with $1/\delta_C$.

In order to choose the proper value of δ_C , we solved a representative subset of 54 trips (see below for the dataset description) using the QP solver in the package CVX-OPT [21]. We obtained a total i^2 -score of $3.070 \cdot 10^8$ in about 11 minutes. Using our DP solver with $N = 90$ steps, we obtained a score of $3.103 \cdot 10^8$ in 15 seconds; with $N = 45$ steps we obtained $3.197 \cdot 10^8$ in about 3 seconds. This experiment empirically shows that our DP based solver is about 2 orders of magnitude faster than solving the quadratic program directly, and provides solutions that are close to optimal. These experiments will guide the choice of the discretization step also for the original stochastic setting where the demand is not known ahead of time.

Robustness Since we will later use supervised learning techniques to learn a policy from a training set, we are interested in measuring how robust is the policy to implementation errors on the actions. The sensitivity plot in Figure 3 is obtained by artificially adding i.i.d. Gaussian noise with variance σ^2 to the optimal action given by the optimal omniscient policy. The plot is averaged over a subset of trips and shows that the performance degrades smoothly as a function of the variance of the implementation error.

Rolling horizon In order to compute the optimal omniscient policy, we need to know the entire future demand $\{d(t), t = 0, \dots, T - 1\}$ ahead of time. Relaxing this assumption, we now assume to know the future demand only for a window of M steps. We then use a rolling horizon policy, where at every step t we replan computing the optimal sequence of actions for the next M steps, taking the first one. In Figure 4 we




Fig. 3: Performance with noise in the implementation of the optimal policy. On the y-axis is the percentage reduction in the i^2 -score with respect to baseline (no capacitor).

show how the performance improves as the length of the rolling horizon window M increases. In particular, notice that if we were able to predict (exactly) the future demand for the next 30 steps (corresponding to 30 seconds), on average we would lose less than 5% over the optimal omniscient policy (that knows the entire future demand ahead of time). This experiment suggests that even a fairly limited probabilistic model that can predict the future demand for a few seconds could provide substantial energy savings, and motivates our search for an MDP model.

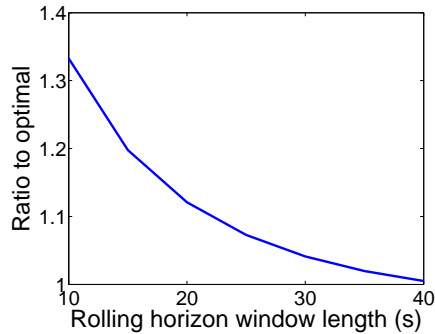


Fig. 4: Performance improvement as the length of the rolling horizon window increases.

5 Probabilistic Planning

An MDP model In general, we cannot know in advance what will be the future demand (load profile), but we can assume the existence of an underlying stochastic model, from which the trips and driving behaviors we observe are sampled from. Specifically, we consider a *state space* $S = \mathcal{F} \times [0, C_{max}]$, where \mathcal{F} is a *feature space*. The idea

is to use a set of features $\mathbf{f} = (f^1, \dots, f^K) \in \mathcal{F}$ and a capacity level $0 \leq c \leq C_{max}$ to represent the state of the electric vehicle at any given time. The features we use are driver ID (and type of vehicle), GPS latitude and longitude, direction, speed, acceleration, altitude, instantaneous demand d , past average demand, time of the day. According to the problem definition, for any state $s = (\mathbf{f}, c) \in \mathcal{S}$, there exists a set of admissible actions $\mathcal{A}(s) = \{\Delta, -c \leq \Delta \leq C_{max} - c\}$, i.e. the admissible changes of the capacitor level that satisfy the constraints of the problem.

Our underlying assumption is that there exist a probabilistic model describing the evolution of the state $P(s_{t+1}|s_0, \dots, s_t, \Delta_0, \dots, \Delta_t)$ as a function of the previous history and the sequence of action $\Delta_0, \dots, \Delta_k$ taken. Further, we assume that

$$\begin{aligned} P(s_{t+1}|s_0, \dots, s_t, \Delta_0, \dots, \Delta_t) = \\ P(\mathbf{f}_{t+1}|\mathbf{f}_0, \dots, \mathbf{f}_t) P(c_{t+1}|c_t, \Delta_t) \end{aligned} \quad (4)$$

that is the evolution of \mathbf{f}_t is independent of c_t and the actions Δ_t taken (equivalently, we assume the driving behavior and road conditions do not depend on the capacitor charge levels). On the other hand, according to the problem description c_{t+1} depends *deterministically* on the past, specifically $c_{t+1} = c_t + \Delta_t$. In this way, $P(c_{t+1}|c_t, \Delta_t) = 1$ if and only if $c_{t+1} = c_t + \Delta_t$.

In this MDP framework, a energy management system is a function mapping histories to a feasible action, i.e. a history-dependent feasible policy [22]. Upon defining an immediate cost $c(\Delta, s, s') = (d + \Delta)^2$ for transitioning from state s to state s' when taking action Δ (equal to the squared current output from the battery), an optimal energy management system can be defined as one minimizing the total expected cost.

A sample-based approach Since the probabilistic model $P(\mathbf{f}_{t+1}|\mathbf{f}_0, \dots, \mathbf{f}_t)$ is unknown, we use a sample-based approach where we leverage a large dataset of commuter trips crowdsourced in the United States and available online [8] in order to learn it from the data. Specifically, we assume that each trip in the dataset corresponds to one particular realization of the underlying stochastic process, e.g. a sampled trajectory of length T_i in the feature space \mathcal{F} . In particular, we project the trip data on the feature space \mathcal{F} , generating a trajectory $\mathcal{T}_i = (\mathbf{f}_0, \mathbf{f}_1, \dots, \mathbf{f}_{T_i-1})$ where for each time step $\mathbf{f}_t \in \mathcal{F}$. Then, for any sequence of actions $(\Delta_0, \dots, \Delta_{T_i-1})$ and initial capacity level c_0 , we can generate the corresponding trajectories in the *state* space $((\mathbf{f}_0, c_0), (\mathbf{f}_1, c_1), \dots, (\mathbf{f}_{T_i-1}, c_{T_i-1}))$ according to (4).

Let $\mathcal{T}_1, \dots, \mathcal{T}_K$ be the sample trajectories available. For any state $s = (\mathbf{f}, c) \in S$ we define the multiset

$$\mathcal{N}(\mathbf{f}) = \bigcup_{i=1}^K \{\mathbf{f}_{t+1}|\mathbf{f}_t \in \mathcal{T}_i, \|\mathbf{f}_t - \mathbf{f}\|_\infty < \epsilon\} \subseteq \mathcal{F}$$

that when $\epsilon = 0$ corresponds to the set of feature vectors that we have observed occurring immediately after \mathbf{f} in the sample trajectories. In practice, since our feature space is continuous we choose $\epsilon > 0$ to discretize the space, so that two feature vectors are considered to be the same if they are close enough (e.g., when a driver comes to an intersection with approximately the same speed, acceleration, etc.). We use k-d trees [23]

to speed up the computation of $\mathcal{N}(\mathbf{f})$ for all observed feature vectors $\mathbf{f} \in \bigcup \mathcal{T}_i$, and in our experiments ϵ is set to one thousandth of the average distance between consecutive feature vectors in the available trajectories. Similarly, we can define a multiset of possible successors in the state space as

$$\mathcal{S}(s) = \mathcal{S}(\mathbf{f}, c) = \bigcup_{c'=0}^{C_{max}} \{(\mathbf{h}, c') | \mathbf{h} \in \mathcal{N}(\mathbf{f})\}$$

Posterior optimal actions We can then define sample-based Dynamic Programming equations as follows

$$V(s) = \min_{\Delta \in \mathcal{A}(s)} \left(\frac{1}{|\mathcal{S}(s)|} \sum_{s' \in \mathcal{S}(s)} V(s') + c(\Delta, s, s') \right)$$

and solve for the ‘‘posterior optimal action’’

$$\Delta^*(s) = \arg \min_{\Delta \in \mathcal{A}(s)} \left(\frac{1}{|\mathcal{S}(s)|} \sum_{s' \in \mathcal{S}(s)} V(s') + c(\Delta, s, s') \right)$$

This approach has the nice theoretical property that the sample-based approximation converges to the true DP equations (for the discretized MDP) in the limit of infinite samples. Similarly, $\Delta^*(s)$ converges to the optimal action as more samples are collected. In contrast, separately optimizing for the single realizations as in [6] (which corresponds to choosing $\mathcal{N}(\mathbf{f}_t) = \mathbf{f}_{t+1}$) doesn’t necessarily converge to the true optimal action as $K \rightarrow \infty$, although it has been shown to work well in practice.

Regressing the optimal policy Using the available sample trajectories $\mathcal{T}_1, \dots, \mathcal{T}_K$, we generate a labeled training set¹ of (state,optimal action) pairs by solving the corresponding sample-based DP equations using value iteration (notice that the ‘‘empirical’’ MDP can have loops, so we cannot solve it in one pass). Specifically, we compute $\Delta^*(\mathbf{f}, c)$ for every $\mathbf{f} \in \mathcal{T}_i$ and for every capacity level $c \in [0, C_{max}]$. We then use supervised learning to learn the relationship between a state $s = (\mathbf{f}, c) = (f^1, \dots, f^K, c) \in S$ and the corresponding optimal action Δ . Notice that the particular structure of the problem allows us to artificially generate N times more data points than what we originally started with. Experimentally, we have seen this to be a crucial improvement in order for the supervised learning algorithm to correctly understand the role of the capacity level c . In particular, we found that generating a dataset just using the optimal sequence of actions \mathbf{a}^* for each trajectory $(\mathbf{f}_0, \mathbf{f}_1, \dots, \mathbf{f}_{T-1})$ is not sufficient to achieve good performance.

The quadratic nature of the cost function gives us further insights on the performance of the supervised learning method used. In particular, the mean-squared error (MSE) is an important error metric in this case, because by reverse triangular inequality

$$\|(\mathbf{a}^* + \mathbf{d}) - (\hat{\mathbf{a}} + \mathbf{d})\|_2 = \|\mathbf{a}^* - \hat{\mathbf{a}}\|_2 \geq \|\hat{\mathbf{a}} + \mathbf{d}\|_2 - \|\mathbf{a}^* + \mathbf{d}\|_2$$

¹ The training dataset will be made available online.

where $\mathbf{a}^* = (\Delta_0^*, \dots, \Delta_{T-1}^*)$ is the optimal sequence of actions, $\hat{\mathbf{a}}$ is the sequence of actions given by regression, and $\mathbf{d} = (d(0), \dots, d(T-1))$ is the demand vector. This gives

$$\|\hat{\mathbf{a}} + \mathbf{d}\|_2 \leq \|\mathbf{a}^* - \hat{\mathbf{a}}\|_2 + \|\mathbf{a}^* + \mathbf{d}\|_2$$

so that $\|\mathbf{a}^* - \hat{\mathbf{a}}\|_2$ bounds the difference in terms of i^2 -score between the optimal sequence of actions \mathbf{a}^* and $\hat{\mathbf{a}}$.

6 Evaluation: the Chargecar Competition

As previously mentioned, we evaluate our method on the publicly available dataset provided by Chargecar [8], a crowdsourcing project open to the public with the goal of making electric vehicles more practical and affordable. Along with the dataset, the Chargecar project provides a simulator to evaluate the performance of power management policies on the trips contained in the dataset. Furthermore, they set up an open algorithmic challenge where the goal of the contest is to design policies that optimize the energy performance of electric vehicles, as measured in terms of the i^2 -score. All the parameters of the model are set as in the competition. In particular, the supercapacitor and the battery provide 50 Watt-hour and 50000 Watt-hour, respectively. Among many other factors, the energy efficiency of multi-battery schemes depends crucially on these parameters. Understanding their interplay with smart energy management policies is one of the goals of the competition, because it would allow us to design better, more efficient electric vehicles.

Dataset The dataset [8] contains a total of 1984 trips (with an average length of 15 minutes), subdivided into 6 separate datasets according to the driver ID. Each one of these dataset is further separated into two subsets: a training and judging set. There are 168 trips in the judging set, accounting for about 8% of the total. Using the trips contained in the training set, we generate a dataset of labeled (state, optimal action) example pairs with the method explained in the previous section. The maximum capacity level is discretized into $N = 45$ discrete steps, and the time step is $\delta_t = 1s$, such that the resulting training dataset contains 75827205 examples. Since the complete training dataset generated with the previously described approach is too big to fit into memory, we divide it according to the driver ID, generating separate training sets for each driver. When these datasets are still too big, we divide them again according to the capacity level feature (selecting entries corresponding to one or more rows of the DP tables). We then learn separate models for each one of these smaller datasets, as shown in figure 5 ².

Supervised learning We used non-parametric exploratory models as there is little or no prior knowledge and possibly highly non-linear interactions. In particular, we use

² A separate default model is trained for previously unseen driver IDs.

bagged decision trees, with the REPTree algorithm as implemented in the Weka package [24] as the base learner. REPTree is a fast regression tree learner that uses information gain as the splitting criterion and reduced-error pruning (with backfitting). Following standard practice, the parameters were set by 5-fold crossvalidation, selecting the model with the best MSE score. We call the resulting policy `DPDecTree`. Using decision trees, we can represent and evaluate the policy efficiently in order to compute the optimal action. In contrast, it can be impractical to compute an action solving an optimization problem based on an estimated future demand, because in a real-world setting it might not be feasible to solve such problems with a high frequency on a car.

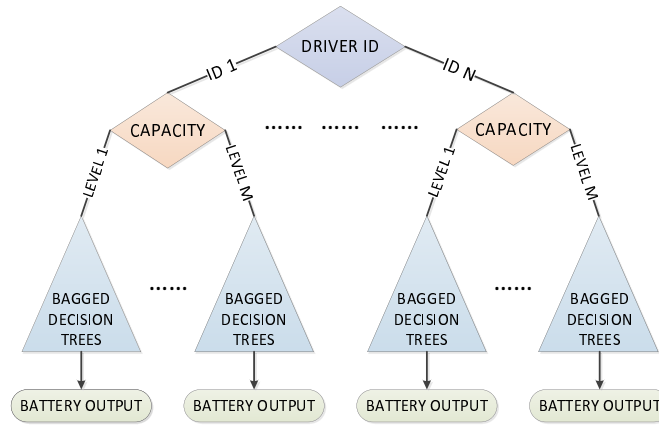


Fig. 5: An overview of our intelligent energy management system.

Evaluation We evaluate the performance of `DPDecTree` on the separate judging set of trips using the simulator. If the action suggested by `DPDecTree` would overcharge the capacitor, we charge it to full capacity; if it overdraws, we discharge it completely. However, this situation is rare and `DPDecTree` gives feasible actions 99.8% of the time when evaluated on the judging set. We compare our solution against `MPL`, the current winning algorithm in the competition at the moment of this paper submission, and to a simple buffer policy. The `MPL` policy is based on a large table of thresholds for the battery output $i_{bc} + i_{bm}$, chosen according to driver ID, speed, demand and GPS coordinates. The naive buffer policy charges the capacitor only from regenerative braking (i.e, when the demand is negative) and when there is an energy demand, it utilizes the capacitor first. We also provide the i^2 -score when the supercapacitor is not used (or is not available) as a baseline.

As can be seen from table 1, `DPDecTree` leads to significant energy savings with respect to the simple buffer policy. Although it is often far from the upper bound represented by the optimal Omniscient policy (except on the *mike* dataset where they differ only by 5%), the good performance of `DPDecTree` suggests that it is very effective at

Table 1: Results on driver specific judging datasets. i^2 -scores are in $10^8 \cdot A^2 \cdot s$.

Dataset	DPDecTree	MPL	Naive Buffer	Baseline	Omniscient
alík	4.233	4.435	7.533	8.424	3.196
arnold	4.090	3.946	8.402	8.894	3.332
mike	3.245	3.290	4.874	5.128	3.083
thor	1.648	1.787	3.931	4.596	1.413
illah	0.333	0.353	0.751	0.856	0.211
gary	2.000	2.146	5.187	5.857	1.261
Total:	15.549	15.957	30.678	33.755	12.496

predicting future energy demands, and that the policy learned from the training examples generalizes to new, previously unseen scenarios. Further, DPDecTree improves over MPL in 5 out of 6 datasets (each one corresponding to a different driver). We believe the problem with the *arnold* dataset is the presence of different driving behaviors in the same dataset that the learning algorithm was not able to separate using the available covariates.

In the bottom row of Table 1 we provide the resulting scores for the entire judging dataset. On average, DPDecTree leads to a 2.5% improvement on the i^2 -score with respect to MPL. According to a one-sided paired t-test, the difference is statistically significant (with P-value 0.0058). The result is significant also according to a Wilcoxon signed-rank test (with P-value 0.00028) [25]. While there is still a significant 20% gap with respect to the optimal omniscient policy, it is not clear how much can still be achieved in the real online setting where the future demand is not known ahead of time. These results suggest that there is a great potential for including routing information (e.g., from a car navigation system) into the problem, since they would bring the policies closer to the omniscient case. However, in this paper we assume that routing information is not available, as in the competition. It would also be very interesting to explore the possibility of improving the policy by learning from new data as it becomes available during the policy evaluation phase, possibly with an incremental learning approach.

7 Conclusions

In this paper we have presented an effective solution to the problem of managing multi-battery systems in electric vehicles. Our novel intelligent energy management system is evaluated on a large dataset of commuter trips crowdsourced in the United States. Our approach is completely data-driven and can be expected to improve as more data is being collected and becomes available.

Our method combines several existing approaches to solve a problem that we model as an MDP with unknown transition probabilities. We use a sample-based approach, where samples are not generated from an analytic model or from a simulator but given as part of a dataset. By observing the empirical transition probabilities of a discretized problem, we solve sample-based dynamic programming equations using value iteration. Thanks to the special structure of the problem and its independent dynamics assumption, we can generate more artificial data points from the samples by exploiting the information contained in the dynamic programming tables. The optimal posterior actions given the observed samples are then combined to form a policy for the original problem. In order to do this, we use supervised machine learning techniques to build a regression model that gives the action as a function of the state of the system. The obtained policy is evaluated on a separate set of real world trips, where it is shown to generalize to situations that were previously unseen in the training set. Our novel system is shown to outperform the leading algorithms that were previously proposed as part of an open algorithmic challenge.

8 Acknowledgements

This work was supported by NSF Expeditions in Computing Grant 0832782.

References

1. C. Gomes. Computational Sustainability Computational Methods for a Sustainable Environment, Economy, and Society. *The Bridge, National Academy of Engineering*, 39(4), 2009.
2. M. Sachenbacher, M. Leucker, A. Artmeier, and J. Haselmayr. Efficient energy-optimal routing for electric vehicles. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.
3. J. Apple, P. Chang, A. Clauson, H. Dixon, H. Fakhoury, M. Ginsberg, E. Keenan, A. Leighton, K. Scavezze, and B. Smith. Green driver: Ai in a microcosm. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.
4. Paul Dille, Matt Duescher, Illah Nourbakhsh, Gregg Podnar, and Joshua Schapiro. Evaluating the urban electric vehicle. In *Carnegie Mellon Technical Report*. February 2010.
5. R. Kötz, S. Müller, M. Bärtschi, B. Schnyder, P. Dietrich, FN Büchi, A. Tsukada, GG Scherer, P. Rodatz, O. Garcia, et al. Supercapacitors for peak-power demand in fuel-cell-driven cars. In *ECS Electro-Chemical Society, 52nd Meeting, San Francisco*, 2001.
6. M. Fox, D. Long, and D. Magazzeni. Automatic construction of efficient multiple battery usage policies. In *Proc. Int. Conf. on Aut. Planning and Scheduling (ICAPS)*, 2011.
7. M. Jongerden, B. Haverkort, H. Bohnenkamp, and J.P. Katoen. Maximizing system lifetime by battery scheduling. In *Dependable Systems And Networks, 2009. DSN'09. IEEE/IFIP International Conference on*, pages 63–72. IEEE, 2009.
8. CreateLab. The chargecar project, January 2012. <http://www.chargecar.org>.
9. W.B. Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. Wiley-Blackwell, 2007.
10. R.S. Sutton and A.G. Barto. *Reinforcement learning: An introduction*, volume 28. Cambridge Univ Press, 1998.
11. L.P. Kaelbling, M.L. Littman, and A.W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

12. P. R. Kumar and Pravin Varaiya. *Stochastic systems: estimation, identification and adaptive control*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1986.
13. D.P. Bertsekas and J.N. Tsitsiklis. *Neuro-dynamic programming*. 1996.
14. J.N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *Automatic Control, IEEE Transactions on*, 42(5):674–690, 1997.
15. G.J. Gordon. Stable function approximation in dynamic programming. In *Machine learning: proceedings of the Twelfth International Conference on Machine Learning, Tahoe City, California, July 9-12, 1995*, page 261. Morgan Kaufmann, 1995.
16. R. Khardon. Learning to take actions. *Machine Learning*, 35(1):57–90, 1999.
17. R. Khardon. Learning action strategies for planning domains. *Artificial Intelligence*, 113(1):125–148, 1999.
18. S.W. Yoon, A. Fern, and R. Givan. Using learned policies in heuristic-search planning. In *Proceedings of the 20th IJCAI*, 2007.
19. S.B. Peterson, J. Apt, and JF Whitacre. Lithium-ion battery cell degradation resulting from realistic vehicle and vehicle-to-grid utilization. *Journal of Power Sources*, 195(8):2385–2392, 2010.
20. Daniel Wong. Energy management optimization in electric vehicles using model predictive control. *Unpublished technical report*, 2011.
21. J. Dahl and L. Vandenberghe. Cvxopt: A python package for convex optimization. In *Proc. Eur. Conf. Op. Res*, 2006.
22. M.L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, Inc., 1994.
23. Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, September 1975.
24. M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I.H. Witten. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.
25. F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.