The Route Not Taken:
Driver-Centric Estimation of Electric Vehicle Range

Peter Ondrúška and **Ingmar Posner**
Mobile Robotics Group, University of Oxford, United Kingdom
{ondruska, ingmar}@robots.ox.ac.uk

Abstract

This paper addresses the challenge of efficiently and accurately predicting an electric vehicle's attainable range. Specifically, our approach accounts for a driver's generalised route preferences to provide up-to-date, personalised information based on estimates of the energy required to reach every possible destination in a map. We frame this task in the context of sequential decision making and show that energy consumption in reaching a particular destination can be formulated as policy evaluation in a Markov Decision Process. In particular, we exploit the properties of the model adopted for predicting likely energy consumption to every possible destination in a realistically sized map in real-time. The policy to be evaluated is learned and, over time, refined using Inverse Reinforcement Learning to provide for a life-long adaptive system. Our approach is evaluated using a publicly available dataset providing real trajectory data of 50 individuals spanning approximately 10,000 miles of travel. We show that by accounting for driver specific route preferences our system significantly reduces the relative error in energy prediction compared to more common, driver-agnostic heuristics such as shortest-path or shortest-time routes.

Introduction

According to recent market forecasts the number of electric vehicles (EVs) on the roads worldwide is set to increase from ca. 150,000 in 2013 to over two million by 2020 (ABI 2013). The adoption of this technology is driven largely by environmental, economic and political factors. However, recent studies have shown that one of the primary impediments to such mass-market adoption is *range anxiety* due to inaccurate in-situ estimates of available vehicle range (Nilsson 2011). As a result, many studies now exist aimed at modelling energy consumption and factors influencing it such as a driver's likely acceleration profile (see, for example, (Karbowski, Pagerit, and Calkins 2012; Oliva, Weihrauch, and Torsten 2013)). Often, these systems provide the user with an indication of attainability for a particular destination queried. However, a significant disadvantage of these approaches is the requirement to manually specify a-priori the desired destination – a task which quickly becomes a nuisance. In addition, even if the desti-

nation were known, the exact route taken may induce variations in energy usage of up to 40% (Minett et al. 2011).

Our work aims to address these shortcomings by providing the driver with a personalised range map, which exhaustively specifies attainability for *every* destination in a realistically sized map in real-time, without the burden on the driver to provide route or destination information (see Fig. 1). Instead of such explicit user-interaction we propose a life-long learning system, which continuously adapts to driver-specific energy needs by learning a route preference model. Attainability estimates are derived by comparing predictions of likely energy consumption with the current state of charge of the EV battery. While our work leverages a commonly employed route preference model, it is the requirement for efficient consideration of all routes to every possible destination in the map, which sets our endeavour apart from the otherwise richly studied area of route prediction. In particular, we show that framing this problem in the context of sequential decision making provides a natural mechanism not only for the incorporation of driver-specific information but also allows for an efficient implementation suitable for real-time deployment. Energy estimates themselves are based on a canonical model of vehicle energy consumption accounting for driving style, route infrastructure and geography.

To the best of our knowledge this is the first work aimed at efficiently providing personalised range maps by accounting for a driver's generalised route preferences only based on observing driving behaviour. Further contributions of our work are

- the formulation of the prediction of expected energy consumption as policy evaluation in a Markov Decision Process,

- an efficient algorithm for providing real-time, personalised predictions of attainability for *all destinations* in a map,

- a system capable of life-long adaptation to user preferences.

We demonstrate that continuously accounting for generalised route preferences over time significantly reduces the error in predicted energy use – and hence improves prediction of attainable range.

Figure 1: *Range maps* personalised for particular drivers as typically provided by our system. The attainability of every possible destination is indicated (shaded area) for the same vehicle location (triangle) and battery state of charge. Our work specifically accounts for the variation of these maps induced by a driver's route preferences by continuously observing trajectories travelled by a specific user.


Figure 2: State-action space of the MDP modelling the road network. States correspond to oriented road segments and actions correspond to possible actions at the end of a road segment.

Problem Formulation

Consider a driver who routinely uses an electric vehicle for transportation in a given area, for which a map is available. Specifically, we consider a map to consist of a road network composed of individual road segments s_i, which are joined at intersections. Such a road network together with associated route infrastructure (e.g. number and location of traffic lights, stop signs, etc) can be readily obtained from community projects such as OpenStreetMap (Haklay and Weber 2008).

The problem of range map generation can be formulated as identifying map destinations which are energetically attainable. We model this situation by considering the destination s_g being attainable if the expected energy required to travel there $E_{\boldsymbol{\theta}}(s_s, s_g)$ from the current location s_s for user-specific driving preferences $\boldsymbol{\theta}$ is less than the current state of charge of the battery, E_{soc}.

Our goal, then, is to estimate the expected energy, $E_{\boldsymbol{\theta}}(s_s, s_g)$ to every destination, s_g. We assume that travel is restricted to the available road network, i.e. that a driver will not drive off-road. The map topology thus gives rise to a *set* of all available trajectories $\mathcal{T}_{s,g} = \{\varsigma_1, \ldots, \varsigma_n\}$, where each ς_i denotes a specific trajectory from s_s to s_g. These trajectories can be substantially different and each one has associated with it a particular energy cost $E(\varsigma_i)$. We further express a driver's preferences over the set of trajectories as a conditioning of the likelihood of a particular trajectory being traversed, $p_{\boldsymbol{\theta}}(\varsigma_i)$, on the set of user-specific parameters $\boldsymbol{\theta}$. In this formulation, therefore, we are interested in computing, for every possible destination s_{goal} in the map, the expected energy consumed in getting there from the current location s_s as

$$E_{\boldsymbol{\theta}}(s_s, s_g) = \sum_{\varsigma_i \in \mathcal{T}_{s,g}} p_{\boldsymbol{\theta}}(\varsigma_i) E(\varsigma_i). \quad (1)$$

While intuitive, this approach to range map computation suffers from two significant drawbacks. Firstly, the set $\mathcal{T}_{s,g}$ can consist of exponentially many trajectories, which renders a direct computation of $E_{\boldsymbol{\theta}}(s_s, s_g)$ infeasible in a real-time context. Secondly, this problem is compounded by the need to compute $E_{\boldsymbol{\theta}}(s_s, s_g)$ for every possible destination s_g in the map. In the remainder of this paper we set out a computationally efficient solution to the task of computing such a driver-specific range map.

Range Prediction via Sequential Decision Making

We consider a model where generalised route preferences of the driver are encoded as a particular policy in a Markov Decision Process (MDP). We assume this policy to optimise an a-priori unknown, user-specific reward structure, which can be learned and, over time, refined from trajectory data using Inverse Reinforcement Learning (IRL). Computation of Equ. 1 for a single destination s_g is framed as an evaluation of this policy in a related MDP, which has its reward structure replaced with one expressing the energy demand of every action based on a canonical model of energy consumption. At the end of this section we present an efficient polynomial-time algorithm which allows computation of $E_{\boldsymbol{\theta}}(s_s, s_g)$ for every destination at once – a property critical to real-time range map computation.

Energy Consumption As MDP

Let an MDP be specified by the tuple $\{\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}\}$ where the set of states is composed of all road segments in a map, $\mathcal{S} = \{s_1, s_2, \ldots, s_N\}$, the set of available actions is

414

Figure 3: Examples of most likely trajectories between the same start and destination induced by three different route preferences. [**left**] minimum time traveled (the trajectory mainly follows motorways), [**middle**] minimum distance traveled and [**right**] custom preference learned from observations.

composed of all possible turns at the end of a road segment, $\mathcal{A} = \{a_1, a_2, \ldots, a_M\}$, the transition model, $\mathcal{P}$, is deterministic and the reward structure, $\mathcal{R}$, associates one reward with each state-action pair. In fact, the specific MDP we consider here encompasses all *oriented* road segments and associated actions as illustrated in Fig. 2. In this case the rewards express the energy cost of choosing a particular action in a particular state, $E(s_i, a_j)$. The deterministic transition model implies that a particular state-action pair leads to a particular next state with certainty, such that $p(s_j \mid s_i, a_k) = 1$. Any trajectory in the road network can then be expressed as a sequence of state-action pairs $\varsigma = \{\{s_0, a_0\}, \{s_1, a_1\}, \ldots, \{s_{n-1}, a_{n-1}\}, \{s_n, a_n\}\}$ where the final state, s_n, is an absorbing state where no further reward is accrued independent of actions taken. In this model every trajectory has an associated total energy cost given by the sum of the individual costs

$$E(\varsigma_i) = \sum_{\{s_t, a_t\} \in \varsigma_i} E(s_t, a_t). \qquad (2)$$

We assume driver route preferences when driving to s_g to result in a stochastic policy π defining a probability distribution over possible actions $p_\pi(a_j \mid s_i)$ at the end of every road segment. The probability of the user taking a particular trajectory can be expressed as the probability of observing the corresponding state-action sequence, such that

$$p_{\boldsymbol{\theta}}(\varsigma_i) = \prod_{\{s_t, a_t, s_{t+1}\} \in \varsigma_i} p(s_{t+1} \mid s_t, a_t) p_\pi(a_t \mid s_t). \qquad (3)$$

Computing $E_{\boldsymbol{\theta}}(s_s, s_g)$ as per Equ. 1 is now equivalent to evaluating the value of policy π at state s_s.

$$E_{\boldsymbol{\theta}}(s_s, s_g) = V_\pi(s_s). \qquad (4)$$

Several efficient methods for policy evaluation exist. Here we employ a method where $V_\pi(s_s)$ can be found as a solution of a system of linear-equations (Sutton and Barto 1998):

$$\begin{aligned} V_\pi(s_i) &= \sum_{a_j} p_\pi(a_j \mid s_i) \Bigg(E(s_i, a_j) \\ &+ \sum_{s_k} p(s_k \mid s_i, a_j) V_\pi(s_k) \Bigg) \end{aligned} \qquad (5)$$

Moreover, this method has the advantage that it produces value $V_\pi(s_i)$ and hence $E_{\boldsymbol{\theta}}(s_i, s_g)$ for all possible starting states s_i at once. Computation of a range map however requires the opposite – $E_{\boldsymbol{\theta}}(s_s, s_i)$ for all possible destinations s_i - an expected energy to reach every destination. In a subsequent section we show how our formulation makes computation of this quantity feasible at no extra cost. First, however, the driver-specific policy π has to be learned.

Driver Model

The policy to be evaluated implicitly induces the driver-specific probability distribution over possible trajectories, $p_{\boldsymbol{\theta}}(\varsigma_i)$ considered in Equ. 1. This becomes apparent when contrasting, for example, the trajectories taken by drivers who minimise travel time or distance (see Fig. 3) with the more complex preferences often exhibited in reality. Following the work of (Ziebart et al. 2008), here we describe how, for an individual driver, both $p_{\boldsymbol{\theta}}(\varsigma_i)$ and $p_\pi(a_j \mid s_i)$ can be derived given a set of trajectories traversed. As in the examples above we implicitly assume a driver to optimise an a-priori unknown cost function which leads to a particular trajectory to a given destination. In particular, we employ the feature-based IRL formulation proposed by (Ziebart et al. 2008) and express this cost as a driver-specific reward structure in an MDP defined over the road network. More specifically, this MDP is identical to the one described in the previous section apart from the reward structure, which is unknown and needs to be recovered. The reward for a given state-action pair, $R_{\boldsymbol{\theta}}(s_i, a_j)$, is assumed to be a weighted linear combination of features, f_{s_i, a_j}, such that

$$R_{\boldsymbol{\theta}}(s_i, a_j) = \boldsymbol{\theta}^\top f_{s_i, a_j}, \qquad (6)$$

where $\boldsymbol{\theta}$ denotes the weight vector. The features express various properties encountered when transitioning from one road segment to another, such as the segment length, the time required to traverse it, the road class (e.g. motorway, dual carriage way, residential, etc.) and number of lanes, angle of turn as well as the number of full stops required due to, for example, the presence of traffic lights or stop signs. The overall reward for a given trajectory, $R_{\boldsymbol{\theta}}(\varsigma_i)$, is computed as the sum of the rewards of the state-action pairs encountered

along it,

$$R_{\boldsymbol{\theta}}(\varsigma_i) = \sum_{\{s_t, a_t\} \in \varsigma_i} R_{\boldsymbol{\theta}}(s_t, a_t). \tag{7}$$

For given weights $\boldsymbol{\theta}$ the probability of a driver taking trajectory ς_i is considered to be proportional to a function exponential in its total reward,

$$p_{\boldsymbol{\theta}}(\varsigma_i) \propto e^{R_{\boldsymbol{\theta}}(\varsigma_i)}. \tag{8}$$

This preference model gives rise to an equivalent stochastic policy in the MDP specifying the probability of taking action a_j in state s_i, which is proportional to the sum of probabilities of trajectories starting with the given action, such that

$$p_{\pi}(a_j \,|\, s_i) \propto \sum_{\varsigma_k : \{s_i, a_j\} \in \varsigma_{k:t=0}} p_{\boldsymbol{\theta}}(\varsigma_k). \tag{9}$$

Maximum Entropy Inverse Reinforcement Learning can then be used to find the weight vector $\boldsymbol{\theta}$ for a given driver based on observed trajectories. This, as well as the computation of the driver specific policy can be carried out efficiently using the forward-backward algorithm described in (Ziebart et al. 2008).

An important attribute of this model is that it does not suffer from the label-bias problem (Lafferty, McCallum, and Pereira 2001). This ultimately leads to a computational shortcut when evaluating this policy as discussed in a subsequent section.

Efficient Range Map Computation

As described above, policy evaluation approaches can be used to compute the expected energy requirement $E_{\boldsymbol{\theta}}(s_s, s_g)$ in reaching a destination s_g from the current location s_s. The generation of a range map, however, in principle requires the application of such a policy evaluation step for every possible destination, which is computationally infeasible. Instead we present a computational shortcut, which allows for the simultaneous evaluation of $E_{\boldsymbol{\theta}}(s_s, s_i)$ for every possible goal s_i in a map.

In particular, instead of the original MDPs used to model trajectory preferences and energy consumption we now consider related MDPs consisting of the same set of states, actions and rewards but with a transition model $\mathcal{P}'$ which encodes the *transpose* of the original road graph. Effectively, therefore, the outcome of every action is reversed, such that $p'(s_k \,|\, s_i, a_j) = p(s_i \,|\, s_k, a_j)$. Note that every trajectory ς_i


Figure 4: A trajectory spanning states and actions in the original road network (a) and the same trajectory in the transposed network (b).

Algorithm 1 Efficient Range Map Computation.

Input:	s_s	position of the car
	E_{soc}	battery charge
	f_{s_i, a_j}	segments features
	$\boldsymbol{\theta}$	routing preferences
	$E(s_i, a_j)$	segments energy consumption
Output:	$drivable(s_i)$	destinations attainability

Compute road segment rewards

1. $R_{\boldsymbol{\theta}}(s_i, a_j) = \boldsymbol{\theta}^{\top} f_{s_i, a_j}$

Transpose the road network

2. $p'(s_k \,|\, s_i, a_j) = p(s_i \,|\, s_k, a_j)$

Compute driver policy π to reach state s_s

3. $Z_{s_s} = 1$

Recursively compute for N iterations:

4. $Z_{s_i a_j} = \sum_{s_k} p'(s_k \,|\, s_i, a_j) e^{R_{\boldsymbol{\theta}}(s_i, a_j)} Z_{s_k}$

5. $Z_{s_i} = \sum_{a_j} Z_{s_i a_j} + 1_{\{s_i = s_s\}}$

6. $p'_{\pi}(a_j \,|\, s_i) = \frac{Z_{s_i a_j}}{Z_{s_i}}$

Solve system of linear equations for $V'_{\pi}(s_i)$

7. $V'_{\pi}(s_i) = \sum_{a_j} p'_{\pi}(a_j \,|\, s_i) \Big(E(s_i, a_j)$
   $+ \sum_{s_k} p'(s_k \,|\, s_i, a_j) V'_{\pi}(s_k) \Big)$

Compute resulting drivability map

8. $E_{\boldsymbol{\theta}}(s_s, s_i) = V'_{\pi}(s_i)$

9. $drivable(s_i) = \begin{cases} \text{true} & \text{if } E_{\boldsymbol{\theta}}(s_s, s_i) \leq E_{soc} \\ \text{false} & \text{otherwise} \end{cases}$

---

in the original MDPs is also feasible in these new MDPs but that it is traversed in reverse (see Fig. 4). Importantly, as the reward structure remains the same, Equ. 8 implies that the probability distribution over trajectories $\mathcal{T}_{s,g}$ from s_s to s_g in the original MDPs is the same as the probability distribution over trajectories $\mathcal{T}'_{g,s}$ from s_g to s_s in the new MDPs. It follows, therefore, that $E_{\boldsymbol{\theta}}(s_s, s_g)$ in the original road network and $E'_{\boldsymbol{\theta}}(s_g, s_s)$ in the transposed version are equivalent,

$$\begin{aligned} E_{\boldsymbol{\theta}}(s_s, s_g) &= \sum_{\varsigma_i \in \mathcal{T}_{s,g}} p_{\boldsymbol{\theta}}(\varsigma_i) E(\varsigma_i) \\ &= \sum_{\varsigma_i \in \mathcal{T}'_{g,s}} p'_{\boldsymbol{\theta}}(\varsigma_i) E(\varsigma_i) = E'_{\boldsymbol{\theta}}(s_g, s_s). \end{aligned} \tag{10}$$

As a result, $E_{\boldsymbol{\theta}}(s_s, s_i)$ can be computed for every s_i at once by first constructing MDPs based on the transposed road network and then solving for the expected energy consumption $E'_{\boldsymbol{\theta}}(s_i, s_s)$ from every starting state s_i to the destination s_s as described previously.

The algorithm is summarised in Algorithm 1. Computation of $E_{\boldsymbol{\theta}}(s_s, s_i)$ consists of three steps. First, we transpose the road network. Then we compute a policy to reach state s_i based on a particular driver's route preferences in the transposed route network using the backward pass described by (Ziebart et al. 2008). The value of this policy for every state is then computed by solving a system of linear equations.

The final range map is obtained by thresholding the energy consumption predictions against the current state of charge of the vehicle battery, E_{soc}.

Vehicle Energy Model

The energy consumed over the course of a trajectory, $E(\varsigma)$, is modelled as the sum of energy costs of state-action pairs along the trajectory and given by Equ. 2. This cost incorporates events specific to the transition, such as potentially stopping at an intersection or changing velocity according to the law of the land, as well as the energy required while traversing road segment s_{t+1} itself. To estimate these values we combine a canonical model of an expected velocity profile with a physical model of the resulting energy demand of the vehicle.

Velocity Profile

We model a driver's chosen velocity, v, and acceleration, $\dot{v}$, as a road segment is traversed by adapting the Intelligent Driver Model (IDM) by (Treiber, Hennecke, and Helbing 2000). The IDM was originally applied in the context of a car-following scenario. It can, however, be employed here by assuming the end of the road segment to be a car moving at velocity v_t equivalent to the desired velocity at the end of the segment. Given the road speed limit v_0 and distance to the end of the segment s, the acceleration of the driver is then given by the ordinary differential equation

$$\dot{v} = \dot{v}_{max}\left(1 - \left(\frac{v}{v_0}\right)^{\delta}\right) + \dot{v}_{min}\left(\frac{(v - v_t)^2}{4s^2}\right). \quad (11)$$

Values of speed limits, positions of stop signs and traffic lights can be extracted from sources such as OpenStreet map (Haklay and Weber 2008). δ denotes a smoothness parameter. The values used are summarised in Table 1.

Physical Model

The energy cost of an action $E(s_i, a_j)$ is obtained by integrating the vehicle's power demand P_{car} during action a_j,

$$E(s_i, a_j) = \int P_{car}(t) dt. \quad (12)$$

The power demand of the car is the sum of the engine power P_{eng} and the auxiliary power demand P_{aux} due to car accessories such as lights, air conditioning, etc, such that

$$P_{car} = P_{eng} + P_{aux}. \quad (13)$$

Figure 5: The forces acting on a car (Guzzella and Sciarretta 2007).


Figure 6: Velocity, acceleration and power demand profile for three adjacent road segments containing a speed limit, a stop sign and a slight elevation.

Constant	Description	Value
$\dot{v}_{max}$	maximum acceleration	$1 m/s^2$
$\dot{v}_{min}$	comfortable deceleration	$-3 m/s^2$
δ	smoothness coefficient	4
m	vehicle mass	$1521 kg$
c_{rr}	rolling friction coefficient	0.015
c_d	aerodynamic drag coefficient	0.25
A_f	frontal area	$2.846 m^2$
ρ	air density	$1.22 kg/m^3$
P_{aux}	auxiliary power consumption	$490W$

Table 1: Physical quantities and values used to model energy consumption.

The engine power demand, P_{eng}, is a function of velocity and acceleration. We assume an efficient system such that

$$P_{eng} = F_{eng} \cdot v, \quad (14)$$

where F_{eng} is the force produced by the engine required to overcome forces acting on the vehicle at a given speed. As depicted in Fig. 5, these forces decompose into components due to acceleration, friction, air resistance and gravitation (Guzzella and Sciarretta 2007).

$$F_{eng} = F_{acc} + F_{friction} + F_{air} + F_g, \quad (15)$$

where

- $F_{acc} = m \cdot \dot{v}$ is the force needed to accelerate the vehicle,
- $F_{friction} = c_{rr} \cdot m \cdot g$ is the rolling resistance,
- $F_{air} = \frac{1}{2} c_d \cdot A_f \cdot \rho \cdot v^2$ is the aerodynamic resistance force,
- $F_g = m \cdot g \cdot sin(\alpha)$ is the hill-climbing force.

The physical constants required to compute these quantities are inspired by those for a Nissan Leaf and are detailed in Table 1. Fig. 6 shows an example of the resulting acceleration, velocity and power demand profile of a vehicle traversing a typical segment of the road network.

Experiments

This section explores the efficacy of the method proposed and provides benchmarking with respect to alternative ap-

proaches. In particular, we set out to demonstrate that integration of driver-specific route preference information provides a significant performance gain in terms of reduced prediction error as compared to more standard, driver agnostic methods assuming, for example, shortest path or traversal time models. As range map generation crucially depends on the accuracy of the underlying estimates of energy consumption we will focus our evaluation on the relative error incurred in $E_{\boldsymbol{\theta}}(s_s, s_g)$.

For evaluation we use Microsoft's GeoLife dataset (see, for example, (Zheng et al. 2008; 2009; Zheng, Xie, and Ma 2010)) containing GPS traces collected from a number of different users. Specifically, our experiments consider 50 different users, each having on average 100 trajectories spanning the urban area of Beijing, China. The road network and infrastructure information for this region covering about $100km^2$ was obtained from the OpenStreetMap project (Haklay and Weber 2008). This road network is very dense and contains a multitude of possible routes between any two places. The resulting MDP contains 80,000 states and 130,000 actions. As the dataset contains GPS traces only, these were segmented into individual trajectories based on time and position information. Next, the Hidden Markov Model-based method described in (Newson and Krumm 2009) was used to match the GPS trajectories to the traversed road segments. Following this preprocessing step, for every user considered 30 trajectories were sampled randomly from the set of trajectories available. The remainder were used for testing. Across all users this resulted in total training and test set sizes of 1500 and 3500 trajectories, respectively.

Prediction Accuracy

For each driver individually a subset of the training set trajectories were used to learn the preference weights, $\boldsymbol{\theta}$. The prediction accuracy of the learned model was then evaluated using all trajectories in the test set. In particular, for every $\{s, g\}$ pair observed in the test set a prediction is made of the expected energy expense incurred by that driver for that trip, $E_{\boldsymbol{\theta}}(s_s, s_g)$. This value is then compared against the energy requirement computed for the actual trajectory taken. As evaluation metric we thus employ the relative prediction error, $\epsilon(\varsigma_i | \boldsymbol{\theta})$, for trajectory ς_i given the learned weight vector, $\boldsymbol{\theta}$, as compared to ground-truth

$$\epsilon(\varsigma_i | \boldsymbol{\theta}) \;=\; 100 \cdot \left| \frac{E_{\boldsymbol{\theta}}(s_s, s_g)}{E(\varsigma_i)} - 1 \right|. \qquad (16)$$

Fig. 7 shows the relative prediction error for 50 users as the amount of training data is increased. This is akin to more trajectory information becoming available over time. The figure indicates that the prediction error decreases rapidly from that obtained using the initial model as the number of training trajectories is increased. As information is added to the model beyond 16 training trajectories improvements are more marginal until, at 28 training trajectories, the overall prediction error reaches ca. 6.5%. In our formulation we compute $E_{\boldsymbol{\theta}}(s_s, s_g)$ as an expectation over the entire set of possible trajectories from start to goal state. The benefit of this approach over considering, for example, only the most
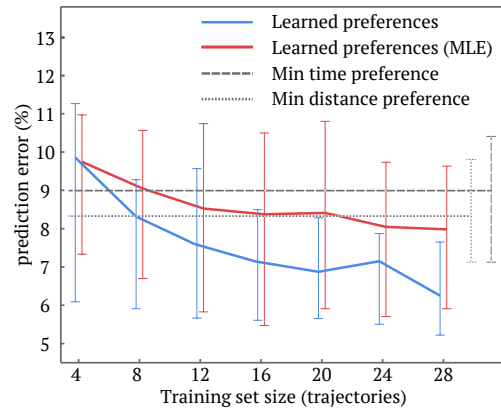

Figure 7: Mean predictive error as the number of training trajectories increases for a variety of models. The approach proposed considers preference learning on a per-driver basis and accounts for the entire distribution over potential trajectories. It achieves the lowest overall prediction error compared to the maximum-likelihood trajectory estimate as well as more common modelling assumptions such as preferences for minimum travel time or distance.

likely trajectory given a driver's preferences is also shown in Fig. 7. The trajectory from start to goal state corresponding to the maximum-likelihood estimate (MLE) is found by computing the shortest path through the road network using the learned segment cost. While the trend of the MLE solution is similar to that found when computing the full expectation in that the prediction error decreases as information is added to the system, the results suggest overall a performance decrease of ca. 1% compared to using the entire distribution. In contrast, our approach computes an expectation over the entire distribution over trajectories.

We further consider two alternative models for path cost: travel time and distance. Results for both of these are also shown in Fig. 7, which indicates that these commonly used estimates result in considerably higher overall prediction error. This is confirmed by a statistical sign test (Gibbons and Chakraborti 2011) at 99% confidence level and suggest that drivers indeed exhibit more complex preferences as to which route is taken.

Timing

The experiments reported here were carried out using code implemented in MATLAB executed on a laptop containing a 2.6GHz i7 processor and 8GB of RAM. With this configuration map updates – i.e. energy usage predictions to every destination for any given starting position – for a map containing 80,000 states were computed at ca. 0.5Hz. Despite running unoptimised code this performance is already sufficient for real-time deployment as map updates can be computed significantly faster than road segments are traversed by the vehicle. We emphasise that this degree of map complexity already encompasses a realistic area of operation. However, as map complexity increases, for example either due to a more elaborate (denser) road network or a greater area covered, map sparsification could be performed

in which adjacent road segments could be merged to obtain similar performance.

Related Works

Not least due to the significant economic implications coupled to a large scale adoption of battery electric vehicles, the prediction of electric vehicle range is an active area of research. Often, considerable effort is expended in generating predictive models for the vehicles' future energy consumption as a proxy for range.

A common way of approaching this problem is to predict a single number – miles to empty – independent of road topology, infrastructure or geography. The approaches by (Ceraolo and Pede 2001), (Chen et al. 2012) and (Oliva, Weihrauch, and Torsten 2013), for example, provide on-line range estimation by extrapolating future energy-use based on that observed in the recent driving history, often of the order of minutes. The accuracy of the prediction at any point in time therefore relies on the implicit assumption that the rate of energy consumption remains relatively constant throughout the journey. This assumption is overly optimistic since the energy consumption of a vehicle varies significantly with factors like acceleration and speed profile, a route's elevation profile, route infrastructure such as traffic lights and stop signs as well as traffic volume (all potentially leading to start-stop behaviour).

A number of works can be considered to overcome these shortcomings. Work by (Karbowski, Pagerit, and Calkins 2012) and (Kim, Lee, and Shin 2013), for example predict the energy consumption of the vehicle along a user-specified trajectory accounting for common factors of influence. In cases when the destination is known but the route is left unspecified a heuristic, such as shortest path or travel time, can be used to obtain an expected route for which energy requirements can then be estimated (see, for example, (Gonzalez et al. 2007)). As our results indicate, however, such predictions are inaccurate because people tend to exhibit more complex route preferences than are captured by these heuristics.

Several works have also considered the more complex prediction problem encountered when neither the destination nor the route are known. (Froehlich and Krumm 2008) and (Joseph, Doshi-Velez, and Roy 2010), for example use anonymised GPS trajectory records of drivers traversing an area to predict the expected trajectory in the near future. Although this method is powerful in some scenarios it requires a good coverage of prior GPS trajectories and can often reliably provide only short-distance predictions.

One way of avoiding altogether the requirement to estimate a destination is to make predictions of energy usage for every possible destination in a map. This results in a map indicating attainable destinations similar to the range maps considered in this work, which provide an intuitive visual reference as to whether a driver's intended destination is reachable or not. (Ferreira, Monteiro, and Afonso 2012) provide a cloud-based solution based on querying an external route planning engine to estimate expected trajectories to nearby destinations. Energy usage estimates are then obtained for a sparse subset of possible destinations by employing trajectory-based models as above. This approach is computationally costly as every destination requires a separate query.

Our work shares the aspiration of providing a driver with a map indicating attainable driving range. Above and beyond prior art, however, our approach accounts for driver-specific route preferences as well route-specific factors by integrating over driver-specific distributions over possible trajectories to every possible destination. Our results show that this leads to more accurate predictions of energy usage over driver-agnostic approaches. Further, and in contrast to (Ferreira, Monteiro, and Afonso 2012), we generate dense maps indicating attainability for *every* possible destination at a speed suitable for on-line deployment.

To the best of our knowledge this is the first work combining route preference modelling with energy usage prediction for electric vehicles. However, our work is closely related to that by (Vogel et al. 2012) who also use IRL to learn route preferences in order to optimise the powertrain controller of a hybrid car, leading to fuel saving of 1.22% over a traditional approach. In contrast to this work, the estimation of an attainable range map poses significant algorithmic challenges – most notably the requirement to efficiently compute estimates for every possible destination and not only short-term behaviour.

Conclusions and Future Work

Previous work in range estimation for electric vehicles has largely ignored driver preferences in terms of route selection. In this work we propose an efficient framework to account for such preferences and show that our approach can reduce the error in predicted energy use to almost one half of that obtained when predicting energy consumption based on common heuristics such as shortest travel time. We leverage an Inverse Reinforcement Learning approach to provide a life-long learning system in a manner entirely transparent to the user – simply by observing the trajectories along which a vehicle travels.

While the results obtained in this work are promising indeed, several avenues are apparent for potential further investigation. Firstly, while an average relative prediction error of 6.5% is achieved using a training set of 28 trajectories per driver we have not as yet investigated the degree to which this error evolves as further information is added to the system. More training trajectories are an obvious candidate for expansion. Furthermore, while the work presented here makes use of a canonical vehicle energy model we hope in the future to be able to use real energy data from a battery electric vehicle. Our approach also does not currently consider second-order effects whereby driving behaviour is influenced by a driver's knowledge of the current state of charge of the battery. In order to explore both of these directions we are currently engaged in obtaining a dataset of real energy usage recorded with a Nissan Leaf.

Acknowledgments

References

2013. Full electric vehicles. Technical Report AN-1564, Allied Business Intelligence (ABI) Research, Intelligent Transportation Systems Technology Analysis.

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

Chen, Y.; Huang, C.; Kuo, Y.; and Wang, S. 2012. Artificial neural network for predictions of vehicle drivable range and period. In *Proceedings of the IEEE International Conference on Vehicular Electronics and Safety*.

Ferreira, J. C.; Monteiro, V.; and Afonso, J. L. 2012. Data mining approach for range prediction of electric vehicle.

Froehlich, J., and Krumm, J. 2008. Route prediction from trip observations. *Society of Automotive Engieers Special Publication (SAE SP)* 2193:53.

Gibbons, J. D., and Chakraborti, S. 2011. *Nonparametric statistical inference*. Springer.

Gonzalez, H.; Han, J.; Li, X.; Myslinska, M.; and Sondag, J. P. 2007. Adaptive fastest path computation on a road network: a traffic mining approach. In *Proceedings of the 33rd international conference on Very large data bases*, 794–805. VLDB Endowment.

Guzzella, L., and Sciarretta, A. 2007. *Vehicle propulsion systems*. Springer.

Haklay, M., and Weber, P. 2008. Openstreetmap: User-generated street maps. *Pervasive Computing, IEEE* 7(4):12–18.

Joseph, J. M.; Doshi-Velez, F.; and Roy, N. 2010. A bayesian nonparametric approach to modeling mobility patterns. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

Karbowski, D.; Pagerit, S.; and Calkins, A. 2012. Energy consumption prediction of a vehicle along a user-specified real-world trip. In *Proceedings from the Electric Vehicle Symposium (EVS)*.

Kim, E.; Lee, J.; and Shin, K. G. 2013. Real-time prediction of battery power requirements for electric vehicles. In *Proceedings of the ACM/IEEE 4th International Conference on Cyber-Physical Systems (ICCPS)*.

Lafferty, J.; McCallum, A.; and Pereira, F. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the ICML*, 336–343.

Minett, C. F.; Salomons, A.; Daamen, W.; Van Arem, B.; and Kuijpers, S. 2011. Eco-routing: comparing the fuel consumption of different routes between an origin and destination using field test speed profiles and synthetic speed profiles. In *Integrated and Sustainable Transportation System (FISTS), 2011 IEEE Forum on*, 32–39. IEEE.

Newson, P., and Krumm, J. 2009. Hidden markov map matching through noise and sparseness. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*, 336–343.

Nilsson, M. 2011. Electric vehicles: The phenomenon of range anxiety. In *Report for the ELVIRE Project (FP7 PROJECT ID: ICT-2009.6.1)*.

Oliva, J.; Weihrauch, C.; and Torsten, B. 2013. A model-based approach for predicting the remaining driving range in electric vehicles.

Sutton, R., and Barto, A. 1998. *Reinforcement learning: An introduction*, volume 1. Cambridge Univ Press.

Treiber, M.; Hennecke, A.; and Helbing, D. 2000. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E 62 (2)* 1805–1824.

Vogel, A.; Ramachandran, D.; Gupta, R.; and Raux, A. 2012. Improving hybrid vehicle fuel efficiency using inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

Zheng, Y.; Li, Q.; Chen, Y.; Xie, X.; and Ma, W. 2008. Understanding mobility based on gps data. In *Proceedings of the 10th International Conference on Ubiquitous Computing*, 312–321.

Zheng, Y.; Zhang, L.; Xie, X.; and Ma, W. 2009. Mining interesting locations and travel sequences from gps trajectories. In *Proceedings of the 18th International Conference on World Wide Web*, 791–800.

Zheng, Y.; Xie, X.; and Ma, W. 2010. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data Engineering Bulletin* 33(2):32–39.

Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 1433–1438.