

Probabilistic Attainability Maps: Efficiently Predicting Driver-Specific Electric Vehicle Range

Peter Ondrůška and Ingmar Posner*

Abstract—This paper concerns the efficient computation of a confidence level with which a particular driver will be able to reach a particular destination given the current state of charge of the battery of an electric vehicle. This probability of attainability is simultaneously computed for all destinations in a realistically sized map while taking into account the driver, the environment, on-board auxiliary systems and the vehicle battery system as potential sources of estimation noise. The model uses a feature-based linear regression framework which allows for a computationally efficient implementation capable of providing real-time updates of the resulting *probabilistic attainability map*. It was deployed on an all-electric Nissan Leaf and evaluated using data from over 140 miles of driving. The system proposed produces results of a quality commensurate with state-of-the-art approaches in terms of prediction accuracy.

I. INTRODUCTION

One of the principal factors impeding the mass market adoption of electric vehicles (EVs) is considered to be *range anxiety* [10]. Various methods exist for predicting remaining miles given the battery charge [2], [11]. These are however not very reliable since the remaining range can vary significantly depending on driving conditions and other influencing factors such as driving style, route preferences and route geography [8], [7], [13]. In light of the limited range of current EVs and a still sparse charging infrastructure the problem of range anxiety is primarily caused by the poor accuracy of instantaneous range predictions provided by the vehicle.

In order to alleviate this problem a number of researchers have recently proposed a generation of range maps indicating the estimated energy requirement to reach a set of destinations and comparing this to the current state of charge of the battery [4], [14]. Commonly these maps require some degree of user interaction or consider only a subset of destinations in the map. The stochastic nature of the underlying processes in terms of energy consumption and state-of-charge estimation are often not considered. (An exception here is [11], which does not provide a map but considers range as a random variable by framing the estimation of the current battery state-of-charge as a recursive estimation problem.)

Recently, [12] have proposed an efficient algorithm to compute personalised attainability estimates for every destination in a map based on route preferences learned from GPS trajectories and a physical model of the likely energy expenditure. Here we build on this approach to provide personalised probabilistic attainability maps which consider

*Peter Ondrůška is a DPhil student and Ingmar Posner is a Professor and a principal investigator at the Mobile Robotics Group, University of Oxford, UK. {ondruska, ingmar}@robots.ox.ac.uk



Fig. 1. Range map prediction installed in a real electric car showing destinations with estimated attainability probability at least 90%.

all sources of uncertainty in the system: the driver, the battery and the environment. The model we propose is entirely data driven and does not make use of an underlying physical model of the battery or the world. Instead we learn a regression model which incorporates current and past energy usage information in order to improve estimates for known trajectories as well as to provide accurate estimates for as yet unknown routes. Due to the fact that energy usage is modelled directly, as opposed to via a physical model of the environment, the approach presented here implicitly captures the influence of local driver behaviour (e.g. acceleration profile) and exogenous factors such as typical traffic flow.

The model is also designed to dovetail with the approach presented in [12] in that a life-long learning paradigm is considered where the learning is done in a manner entirely transparent to the driver and the same efficient algorithm can be used to achieve map updates in real-time for maps of realistic size. To the best of our knowledge this is the first work to propose a computationally efficient, driver-specific and probabilistic treatment of destination attainability. A system overview is presented in Fig. 5. Our approach is implemented on Nissan Leaf (see Fig. 1) and evaluated using real data recorded over 140 miles of driving.

II. PREDICTING ENERGY CONSUMPTION

Consider an electric vehicle which is routinely deployed 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 joint at intersections and set of actions a_i corresponding to possible turns at the end of each road segment, Fig. 2.

Given the current location of the vehicle, s_{start} , and the

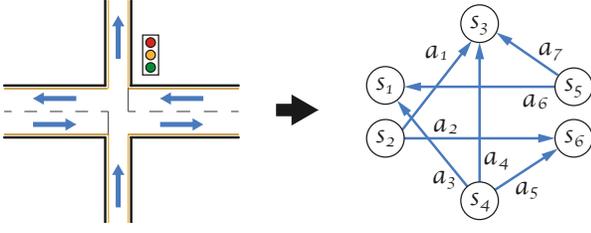


Fig. 2. State-action space of the road network. States correspond with oriented road segments and actions correspond with possible actions at the end of a road segment.

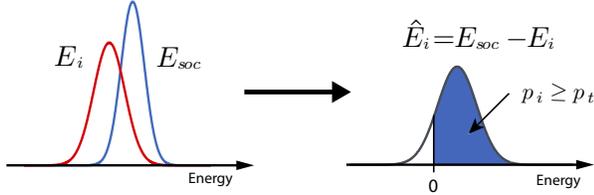


Fig. 3. A statistical interpretation of destination attainability p_i given the distributions of required energy E_i and available energy E_{soc} .

current state of charge of the vehicle battery, E_{soc} , the goal of this work is to provide a framework which efficiently computes an accurate estimate of the probability that a particular destination can be reached for *every* destination in the map. In particular, we estimate a probability p_i that energy E_i required to travel to destination s_i is at most the available energy E_{soc}

$$p_i = p(E_{soc} \geq E_i) \quad (1)$$

and consider this destination attainable if this probability is at least commensurate with a user-specified confidence threshold p_t

$$\Gamma_i = \begin{cases} 1 & \text{if } p_i \geq p_t \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Equivalently to Eq. 1, the value of p_i can be expressed in terms of a transformed random variable $\hat{E}_i = E_{soc} - E_i$ (see Fig. 3) such that

$$p_i = p(\hat{E}_i \geq 0), \quad (3)$$

which leads to an intuitive formulation in terms of the cumulative distribution function of \hat{E}_i , $P_{\hat{E}_i}(x)$, as

$$p_i = 1 - P_{\hat{E}_i}(0). \quad (4)$$

If we assume that the original quantities involved are normally distributed (i.e. $E_i \sim \mathcal{N}(\mu_i, \sigma_i)$ and $E_{soc} \sim \mathcal{N}(\mu_{soc}, \sigma_{soc})$) it follows that \hat{E}_i is also normally distributed such that

$$\hat{E}_i \sim \mathcal{N}(\hat{\mu}_i, \hat{\sigma}_i), \quad (5)$$

$$\hat{\mu}_i = \mu_{soc} - \mu_i, \quad (6)$$

$$\hat{\sigma}_i^2 = \sigma_{soc}^2 + \sigma_i^2. \quad (7)$$

Fig. 4 demonstrates attainability of destinations for different levels of p_t . Eq. 4 can then be expressed in terms of

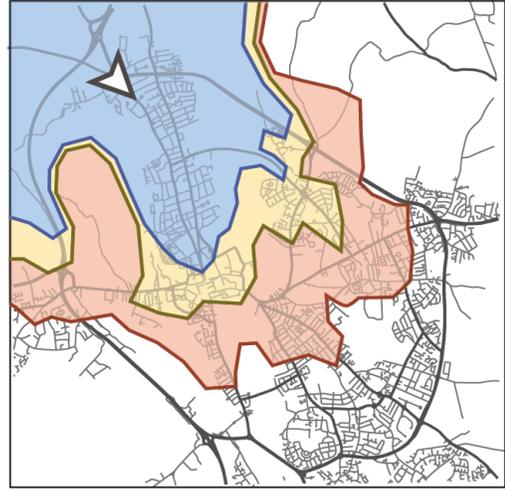


Fig. 4. Attainable range with 90% (blue), 50% (yellow) and 10% (red) probability.

the cumulative distribution function of the standard normal distribution such that

$$p_i = 1 - \Phi\left(-\frac{\hat{\mu}_i}{\hat{\sigma}_i}\right). \quad (8)$$

Note that this formulation caters for the state of charge estimate to be a random variable with associated measurement noise, which can be considerable [3]. While it is conceivable that these parameters are also learnt as part of this model, in the following we will use typical values for μ_{soc} and σ_{soc} as discovered in prior art.

In this context, therefore, the challenge of estimating the chance of reaching a particular destination with a given state of charge consists of accurately and efficiently estimating the parameters μ_i and σ_i , which characterise the likely energy use for any given road segment. Our modelling framework builds on our prior work [12] which uses a feature-based approach to model driver behaviour in the form of route preferences which was first introduced in [15]. We next give a brief summary of this model before extending it to estimating energy values for individual road segments.

A. Driver Model

A driver's route preferences can considerably skew the energy required to reach a particular destination [8], [7]. To account for this variation [12] learns personalised route preferences directly from observed driving behaviour. In particular, consider a trajectory, ς , which consist of an ordered set of road segments. The probability of traversing a particular trajectory is characterised via a reward, R_ς , such that

$$p(\varsigma) \propto e^{R_\varsigma}. \quad (9)$$

The reward for the trajectory is derived from rewards for individual state-action pairs, $\{s, a\}$, along the trajectory, where states represent road segments and actions represent actions potentially taken at the end of a segment (e.g. turning). These individual rewards are modelled as a linear

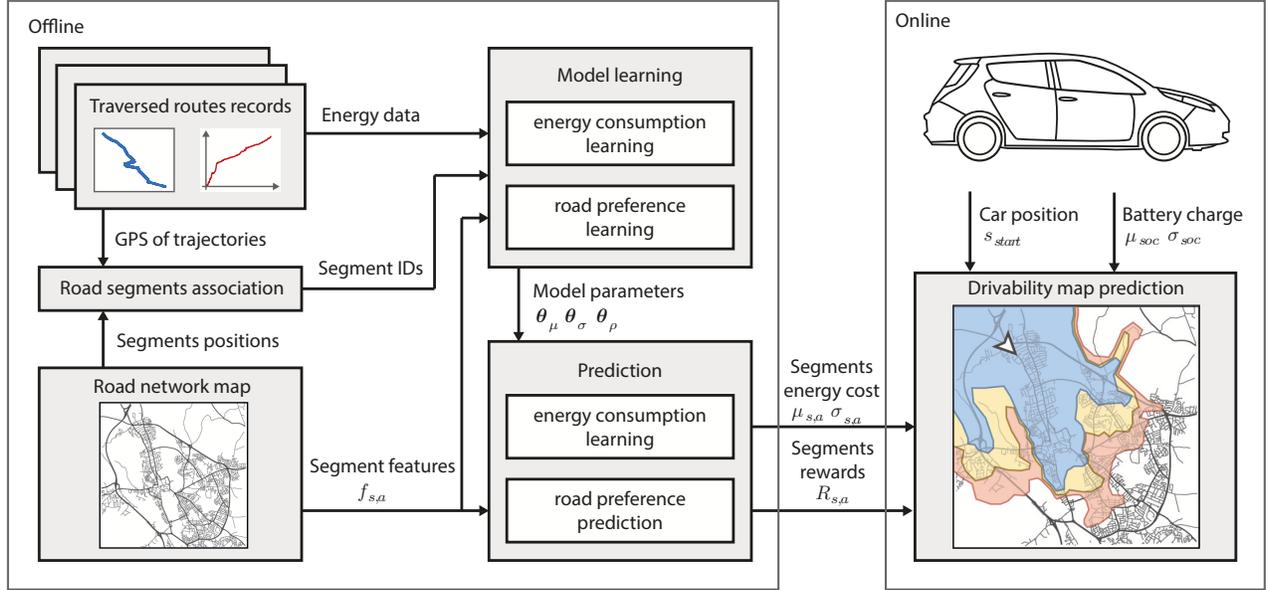


Fig. 5. System architecture overview.

combination of feature values parameterised by a driver-specific weight vector θ_ρ ,

$$R_\zeta = \sum_{\{s,a\} \in \zeta} R_{s,a} \quad (10)$$

$$R_{s,a} = \theta_\rho^\top \cdot f_{s,a} \quad (11)$$

where $f_{s,a}$ are features of the particular road segment and θ_ρ is learnt from data using inverse reinforcement learning [15]. Features of road segments can be extracted from community projects like OpenStreetMap [6]. In this work they consist of the road class (highway, 1st, 2nd, 3rd, residential, service, etc.), the segment length, the amount of turn at the end of the segment (straight, turn left, turn right), whether a stop is likely at the end of the segment (e.g. due to road priorities or the presence of a traffic light or stop sign) as well as the total elevation increase and decrease along the segment obtained from Google Maps Elevation Service.

B. Energy Model

Here we outline the framework used to model energy consumption on route to a particular destination. As in [12] we consider the energy required to reach a particular destination to be computed with respect to a distribution over possible trajectories to that destination. In contrast to [12], however, we consider the energy used for a particular trajectory, E_ζ to be a random variable with distribution parameters derived from real data. E_i is therefore a mixture of distributions describing the energy consumed along possible trajectories, E_ζ , with the mixing coefficients provided by the probability of a driver taking that particular trajectory

$$E_i = \sum_{\zeta} p(\zeta) E_\zeta. \quad (12)$$

As there is potentially a very large number of trajectories connecting two destinations, an exact evaluation of equ. 12 is often infeasible. For computational efficiency and similar to [12], we consider only the set of all trajectories consisting of less than N segments for a large value of N and also re-approximate E_i with a normal distribution specified by [5]

$$\mu_i = \sum_{\zeta} p(\zeta) \mu_\zeta, \quad (13)$$

$$\sigma_i^2 = \sum_{\zeta} p(\zeta) ((\mu_\zeta - \mu_i)^2 + \sigma_\zeta^2). \quad (14)$$

We model the energy of a trajectory as a sum over its traversed road segments

$$E_\zeta = \sum_{\{s,a\} \in \zeta} E_{s,a} \quad (15)$$

Assuming energy expenditure on all road segments is normally distributed and independent from each other implies E_ζ is also normally distributed with given mean and variance:

$$E_{s,a} \sim \mathcal{N}(\mu_{s,a}, \sigma_{s,a}) \quad (16)$$

$$E_\zeta \sim \mathcal{N}(\mu_\zeta, \sigma_\zeta), \quad (17)$$

$$\mu_\zeta = \sum_{\{s,a\} \in \zeta} \mu_{s,a}, \quad (18)$$

$$\sigma_\zeta^2 = \sum_{\{s,a\} \in \zeta} \sigma_{s,a}^2. \quad (19)$$

The parameters $\mu_{s,a}, \sigma_{s,a}^2$ are modelled using a feature-based linear model similar to the one employed to model segment

rewards,

$$\mu_{s,a} = \boldsymbol{\theta}_\mu^\top \cdot f_{s,a}, \quad (20)$$

$$\sigma_{s,a}^2 = \boldsymbol{\theta}_\sigma^\top \cdot f_{s,a}, \quad (21)$$

where $f_{s,a}$ again denotes a feature vector associated with every state-action pair and $\boldsymbol{\theta}_\mu$ and $\boldsymbol{\theta}_\sigma$ denote driver specific weights parameterising the energy consumption. The weights are learned by maximising the likelihood of observed energy consumption (Eq. 16) over the lifetime of the system by constrained optimisation subject to $\boldsymbol{\theta}_\sigma^\top \cdot f_{s,a} \geq 0 \forall s,a$ [1]. To accommodate for different energy demand caused by different accessory settings and selected ECO mode ¹ we learn the weights for specific settings of variables thus we obtain different set of weights for different settings of the car. During prediction the appropriate weights are selected to reflect the current car settings.

III. IMPLEMENTATION

In this section we describe the efficient computation of the parameters μ_i and σ_i as required for the generation of probabilistic attainability maps. Consider an arbitrary function $G_{s,a}$ factorizing over state-action space. In essence, the generation of an attainability map requires the computation of an expectation g_i of $G_{s,a}$ over the distribution over all possible trajectories ς from a starting position to a particular destination s_i

$$g_i = \sum_{\varsigma} p(\varsigma) \sum_{\{s,a\} \in \varsigma} G_{s,a}. \quad (22)$$

In [12] we describe an efficient polynomial-time algorithm, $\mathcal{Q}(\cdot)$, which, given a reward structure R , computes this expectation from a starting position s_{start} to every possible destination s_i in the map all at once, such that

$$\{g_i\} \leftarrow \mathcal{Q}(s_{start}, R_{s,a}, G_{s,a}), \quad (23)$$

where $\{g_i\}$ denotes the set of expectations computed for all destinations. In [12] this algorithm was employed to evaluate eqs. 13 and 18 to obtain the expected mean energy usage, μ_i , by letting

$$G_{s,a} = \mu_{s,a}. \quad (24)$$

Here we extend this approach by also computing the expected variance σ_i^2 . To do this we first execute $\mathcal{Q}(\cdot)$ as above to obtain the mean energy usage for every destination. Consider now a road segment s_k , which is reached by performing action a_j at the end of road segment s_i

$$\bar{\sigma}_{s_i,a_j}^2 = \sigma_{s_i,a_j}^2 + (\mu_{s_i,a_j} + \mu_i - \mu_k)^2 \quad (25)$$

Substituting $\bar{\sigma}_{s_i,a_j}^2$ for $G_{s,a}$ as input to algorithm \mathcal{Q} then allows for the efficient computation of σ_i^2 as output of algorithm \mathcal{Q} . The entire algorithm is summarised in Algorithm 1.

¹ECO mode reduces car acceleration, energy provided to the air conditioning system and increases regenerative braking in order to extend the drivable range.

Algorithm 1 Probabilistic Attainability Map Computation.

Input: s_{start} position of the car
 μ_{soc}, σ_{soc} battery charge
 $R_{s,a}$ segments routing preferences
 $\mu_{s,a}, \sigma_{s,a}$ segments energy consumption
 p_t confidence threshold
Output: Γ_i attainability of destinations

Compute predictive mean

1. $\mu_i = \mathcal{Q}(s_{start}, R_{s,a}, \mu_{s,a})$

Compute predictive variance

2. $\bar{\sigma}_{s_i,a_j}^2 = \sigma_{s_i,a_j}^2 + (\mu_{s_i,a_j} + \mu_i - \mu_k)^2$
3. $\sigma_i^2 = \mathcal{Q}(s_{start}, R_{s,a}, \bar{\sigma}_{s_i,a_j}^2)$

Compute resulting drivability map

4. $\hat{\mu}_i = \mu_{soc} - \mu_i$
5. $\hat{\sigma}_i^2 = \sigma_{soc}^2 + \sigma_i^2$
6. $p_i = 1 - \Phi\left(-\frac{\hat{\mu}_i}{\hat{\sigma}_i}\right)$
7. $\Gamma_i = \begin{cases} 1 & \text{if } p_i \geq p_t \\ 0 & \text{otherwise} \end{cases}$



Fig. 6. The RobotCar (Modified Nissan Leaf) platform used for data gathering.

The algorithm is implemented as a part of a more complex system providing the user real-time updates of attainable range (see Fig. 5). First, in the preprocessing part, records of traversed routes together with extracted road segment features $f_{s,a}$ are used to learn model parameters of energy consumption $\boldsymbol{\theta}_\mu, \boldsymbol{\theta}_\sigma$ and route preferences $\boldsymbol{\theta}_\rho$. These are subsequently used to predict the energy consumption $\mu_{s,a}$ $\sigma_{s,a}$ and user preference cost $R_{s,a}$ for all road segments of the map. Next, in the on-line part, these are taken as an input of the algorithm together with the car position s_{start} and battery charge μ_{soc}, σ_{soc} to produce the attainability of all destinations for a specified confidence threshold p_t . This information is then presented to the user in the form of a visual attainability map.

Our implementation of the algorithm uses a GPU-accelerated computation carried out on a standard MacBook Pro capable of performing prediction every 5 seconds for a map in radius of 100 miles surrounding the car. This is sufficient for a real-time application given the car maximum drivable range, it's position and battery charge rate of change.

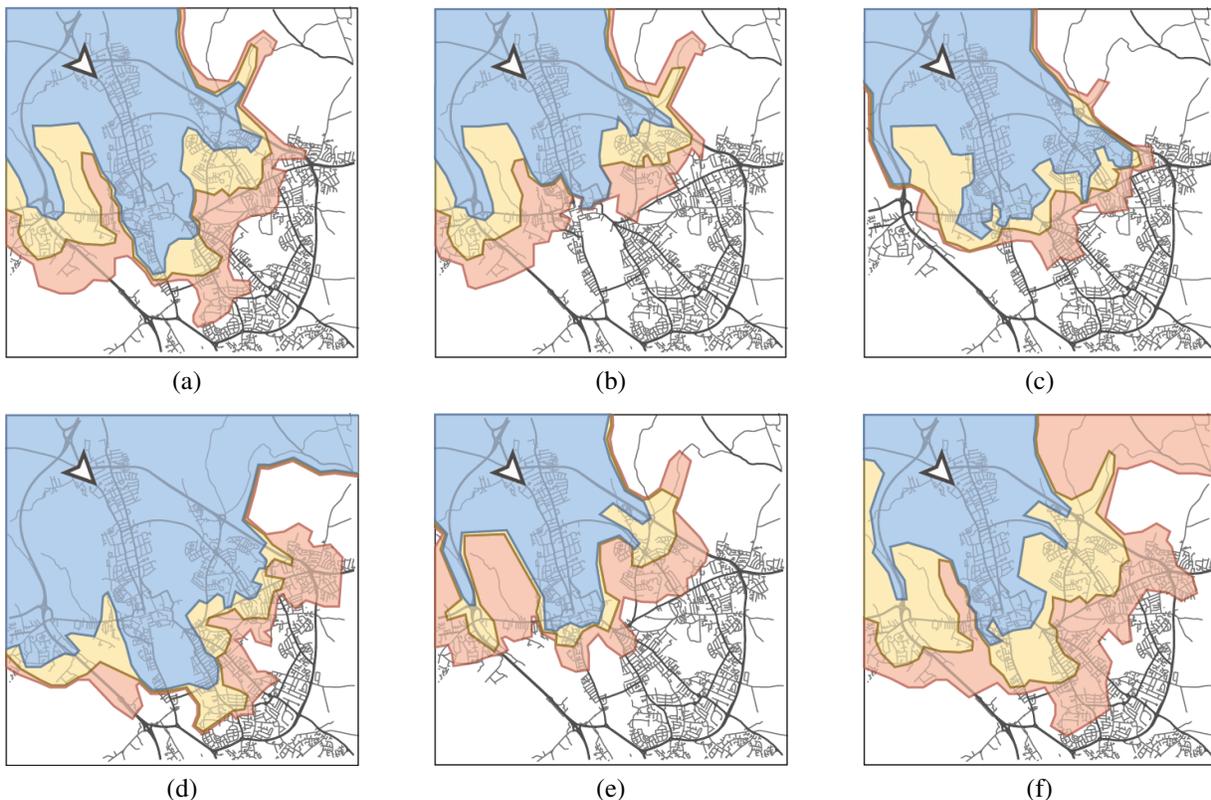


Fig. 7. Examples of driver specific probabilistic attainability maps showing 90% (blue), 50% (yellow) and 10% (red) confidence bounds for various route preferences, driving styles and levels of battery charge estimation noise. Top row: [a] minimum-time route preference, [b] motorway route preference and [c] motorway avoidance route preference. Bottom row: minimum-time route preference with [d] ECO mode ON, [e] air conditioning ON and [f] 10% uncertainty in battery charge.

IV. EXPERIMENTS

Our approach is evaluated using data gathered over 140 miles of driving using an all-electric Nissan Leaf (Fig. 6). The vehicle has been modified to record positioning information provided by a GPS unit. Energy usage information for various vehicle components and state of the battery charge are recorded as reported by the CAN-bus. The corresponding road-network information (including values for road segment features) required for attainability map computation was extracted from OpenStreetMap [6] for a surrounding area in radius of 100 miles. For data association we matched GPS trajectories to sequences of traversed road segments using a Hidden Markov Model as described in [9]. In order to perform a meaningful quantitative evaluation a single route was chosen (covering a variety of road types) and traversed ten times by the same driver. Fig. 7 gives an indication of the types of maps produced by our model for a range of scenarios. As the contribution of this work is a probabilistic model of energy consumption, for these plots the driver’s route preferences were set by manually specifying θ_ρ . Parameters for the energy model were learned from data. In particular, the figure shows the attainability map for a given vehicle location and remaining battery charge $\mu_{soc} = 1.2kWh$ and $\sigma_{soc} = \epsilon_{soc} \cdot \mu_{soc}$ for an RMS state-of-charge estimation precision ϵ_{soc} which was, apart from Fig. 7(f), set to $\epsilon_{soc} = 3\%$ as indicated in [3].

#	energy used [Wh]	energy predicted [Wh]	μ error [%]	σ error [%]
1	280.69 \pm 27.17	291.52 \pm 24.28	3.86	10.63
2	102.20 \pm 27.48	102.41 \pm 11.81	0.21	57.00
3	132.99 \pm 13.49	132.60 \pm 27.90	0.29	106.80
4	230.53 \pm 27.35	219.42 \pm 19.76	4.82	27.74
5	368.20 \pm 56.05	403.36 \pm 43.02	9.54	23.23
6	362.58 \pm 50.32	313.76 \pm 66.89	13.46	32.93
7	244.38 \pm 19.23	259.08 \pm 21.90	6.01	13.91
8	464.24 \pm 24.16	433.97 \pm 29.85	6.51	23.54

TABLE I

PREDICTED MEAN, μ , AND STANDARD DEVIATION, σ , COMPARED TO THE EMPIRICALLY OBSERVED DISTRIBUTIONS FOR EIGHT INDEPENDENT TRAJECTORIES.

The evolution of the maps is rather intuitive. Fig. 7(a)-(c) demonstrate the effect of different driving preferences on the attainability maps. While a detailed interpretation is heavily dependent on the particular road network considered, in this case a minimum-time route preference allows to travel furthest, while both motorway and non-motorway preferences are more restrictive. Fig. 7(d) shows an extended driving range when parameters $\theta_\mu, \theta_\sigma$ were trained while the ECO mode was ON, thereby limiting engine power. Similarly, Figure 7(e) shows a reduced range when the air conditioning is turned ON. Finally, increasing uncertainty in the battery state-of-charge estimate leads to increased width

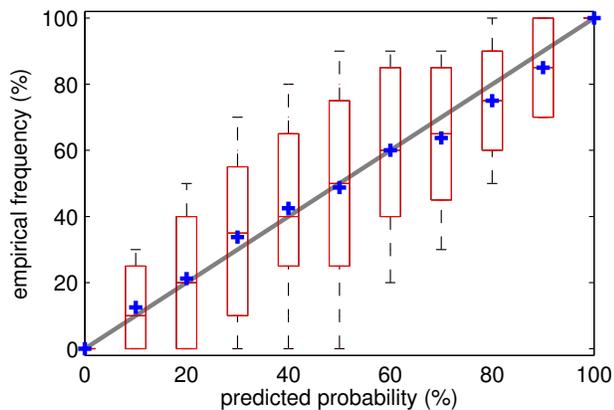


Fig. 8. Empirical proportion of successful attempts to reach destination as a function of predicted attainability probability. Optimal response is linear - marked grey.

of the confidence bands reflecting increased uncertainty in the attainability of destinations. In order to quantitatively evaluate the system we split the dataset into eight non-overlapping trajectories, each of which was traversed ten times. Leave-one out cross-validation was used to provide energy prediction results for each of these trajectories. For these experiments we consider the battery state-of-charge estimate to be noise free. Table I compares the predictions obtained from our system with the empirically measured values. On average our method is able to estimate μ with an average error of 6%, which is commensurate with state-of-the-art approaches such as [8].

To assess the accuracy of our system's predictions of attainability we compare for each one of the eight test trajectories the predicted probability of attainability with the empirically observed one. In particular, for each test trajectory, we vary the available battery energy to control the probability of traversing it to its end. We then compare this to the empirical frequency of successful traversals which require up to that amount of energy. The results of these experiments are shown in Fig. 8, which plots the empirical frequencies against the ones obtained by our model. If the system was perfectly calibrated (i.e. the predictions are exactly right) the relationship would be perfectly linear (grey line). The figure suggests that, while noise exists due to prediction errors, on average the predictions made indeed correspond to those observed.

V. CONCLUSIONS

This paper introduces an efficient approach to computing probabilistic attainability maps for electric vehicles. It considers energy consumption to be a random variable and leverages a feature-based linear regression framework to model the distribution parameters associated with it. Due to the fact that energy usage is modelled directly, as opposed to via a physical model of the environment, the approach presented here implicitly captures the influence of local driver behaviour (e.g. acceleration profile) and exogenous factors such as typical traffic flow. Noise in the vehicle

battery state-of-charge estimate is explicitly accounted for. The simplicity of the model allows it to dovetail with an efficient algorithm to model a driver's route preferences. It therefore achieves a wide coverage of sources of variation while providing for real-time update rates for realistically sized maps. Despite the simplicity of the model (and a number of distribution assumptions) our method provides prediction accuracies of a quality commensurate with state-of-the-art models while, in addition, providing confidence bounds which are well calibrated vis-a-vis observed data.

ACKNOWLEDGMENT

The authors would like to gratefully acknowledge support of this work by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/K034472/1 and via the DTA scheme. We would also like to thank Chris Prahacs, Chi Hay Tong and Dan Withers for generously donating time and energy for data gathering.

REFERENCES

- [1] Richard H Byrd, Jean Charles Gilbert, and Jorge Nocedal. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, 89(1):149–185, 2000.
- [2] M. Ceraolo and G. Pede. Techniques for estimating the residual range of an electric vehicle. *IEEE Transactions on Vehicular Technology*, 50(1), 2001.
- [3] Mohammad Charkhgard and Mohammad Farrokhi. State-of-charge estimation for lithium-ion batteries using neural networks and ekf. *Industrial Electronics, IEEE Transactions on*, 57(12):4178–4187, 2010.
- [4] João C Ferreira, Vítor Monteiro, and João L Afonso. Data mining approach for range prediction of electric vehicle. In *Conference on Future Automotive Technology - Focus Electromobility*, 2012.
- [5] Sylvia Frühwirth-Schnatter. *Finite mixture and Markov switching models*. 2006.
- [6] M. Haklay and P. Weber. Openstreetmap: User-generated street maps. *Pervasive Computing, IEEE*, 7(4):12–18, 2008.
- [7] T. Kono, T. Fushiki, K. Asada, and K. Nakano. Fuel consumption analysis and prediction model for eco route search. In *Proceedings of the 15th World Congress on Intelligent Transport Systems and ITS Americas 2008 Annual Meeting*, 2008.
- [8] Claire F Minett, AM Salomons, Winnie Daamen, Bart Van Arem, and Sjon Kuijpers. 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*, pages 32–39. IEEE, 2011.
- [9] P. Newson and J. Krumm. Hidden markov map matching through noise and sparseness. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*, pages 336–343, 2009.
- [10] M. Nilsson. Electric vehicles: The phenomenon of range anxiety. In *Report for the ELVIRE Project (FP7 PROJECT ID: ICT-2009.6.1)*, 2011.
- [11] J. Oliva, C. Wehrauch, and B. Torsten. A model-based approach for predicting the remaining driving range in electric vehicles. In *Annual Conference of the Prognostics and Health Management Society*, 2013.
- [12] P. Ondruska and I. Posner. The route not taken: Driver-centric estimation of electric vehicle range. In *24. International Conference on Automated Planning and Scheduling (ICAPS)*, 2014.
- [13] Hai Yu, Finn Tseng, and Ryan McGee. Driving pattern identification for ev range estimation. In *Electric Vehicle Conference (IEVC), 2012 IEEE International*, pages 1–7. IEEE, 2012.
- [14] Yuhe Zhang, Wenjia Wang, Yuichi Kobayashi, and Keisuke Shirai. Remaining driving range estimation of electric vehicle. In *Electric Vehicle Conference (IEVC), 2012 IEEE International*, pages 1–7. IEEE, 2012.
- [15] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1433–1438, 2008.