

Enabling Intelligent Energy Management for Robots using Publicly Available Maps

Oliver Bartlett¹, Corina Gurau¹, Letizia Marchegiani¹ and Ingmar Posner¹

Abstract—Energy consumption represents one of the most basic constraints for mobile robot autonomy. We propose a new framework to predict energy consumption using information extracted from publicly available maps. This method avoids having to model internal robot configurations, which are often unavailable, while still providing invaluable predictions for both explored and unexplored trajectories. Our approach uses a heteroscedastic Gaussian Process to model the power consumption, which explicitly accounts for variance due to exogenous latent factors such as traffic and weather conditions. We evaluate our framework on 30km of data collected from a city centre environment with a mobile robot travelling on pedestrian walkways. Results across five different test routes show an average difference between predicted and measured power consumption of 3.3%, leading to an average error of 6.6% on predictions of energy consumption. The distinct advantage of our model is our ability to predict measurement variance. The variance predictions improved by 84.3% over a benchmark.

I. INTRODUCTION

Robots are limited by heavy and expensive batteries, which makes energy efficiency a key constraint on robot performance. Thus, modelling and managing energy consumption is of vital importance to predict the lifetime and range of autonomous platforms. In this paper we propose a probabilistic, data-driven approach to estimating the energy consumption of a mobile robot on a set of trajectories, whether they have been traversed or not. In particular, we treat the robot as a *black box*, thereby removing our reliance on often unavailable system characteristics. Rather than evaluating the energy consumption due to the robot’s locomotion and sensor hardware as well as computation costs, we measure consumption directly on the routes traversed and attempt to utilise features derived from publicly available maps to extrapolate to energy consumption on previously unseen routes. Figure 1 shows a typical output of our approach. In an urban environment our model predicts the energy consumption for a given set of path segments. Crucially, this information can be used by energy management systems or planning algorithms to forecast the attainability of a new goal given a particular trajectory.

Our approach is to frame this as a regression problem whereby *power* consumption is regressed on a feature representation of individual path segments. We elect to model *power* rather than energy consumption as it increases the robustness of our approach in two important ways: firstly, power estimates are more robust to localisation errors. Consider a situation where the robot momentarily localises on

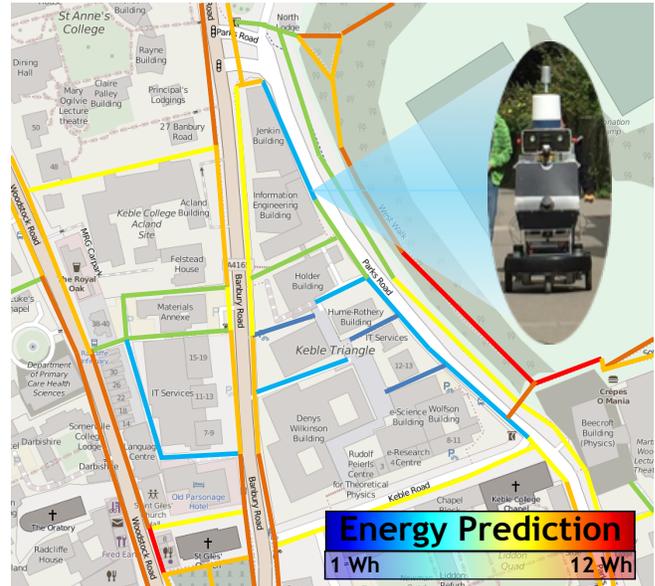


Fig. 1. Example of a typical environment traversed by the robot. The output of our model shows the predicted energy consumption for each segments of our environment, whether visited by our robot or not.

an incorrect path segment. A small energy measurement on that segment will represent a significant outlier, whereas a power reading will still be within a sensible range as it accounts for the amount of time associated with the traversal of a particular segment. For the same reason, power estimates are robust to unpredictable velocity profiles such as caused by traffic or waiting at a pedestrian crossing. We will demonstrate that our power consumption estimates can be converted into predictions of energy consumption by way of a simple velocity profile as provided commonly by a trajectory planner.

While power estimates are considerably more robust to errors in the time spent on a particular path segment, there nevertheless exists a substantial amount of variation due to unknown factors such as pedestrian traffic, weather conditions and different operator behaviours. A crowded environment, for example, may require the robot to manoeuvre to avoid people, thus increasing the variance in power consumption. An uncrowded and open pavement, on the other hand, will likely be traversed similarly each time resulting in a comparatively low variance. To account for this variability, we use the *most likely heteroscedastic Gaussian Process regression* framework proposed in [1], which combines two separate Gaussian Processes (GPs) in order to explicitly accommodate non-uniform measurement noise.

¹Authors are members of Mobile Robotics Group, University of Oxford, United Kingdom, {bartlett, corina, letizia, ingmar}@robots.ox.ac.uk

To the best of our knowledge, this is the first work on predicting the energy consumption of a mobile robot using only high-level, publicly available path features – thereby enabling predictions on unseen trajectories.

The remainder of the paper is structured as follows: Section II investigates the existing literature, in Section III we describe the theory of our approach, Sections IV and V describe our experiments and results using 30km of data collected with a pedestrian robot. We further investigate the predictive power of the individual features within the model.

II. RELATED WORK

Many different methods have been proposed to assess and reduce energy and power consumption in robotic platforms. Motion planning and optimisation of a robot’s speed profile are investigated in [2], [3], [4]. The impact that processing and sensing have on power and energy consumption is analysed in [5], [6], [7]. The possibility of reducing energy consumption by systematically turning off the main localisation system is proposed in [8]. Generally, these approaches rely on the assumption that the power estimates are accurate. This work aims to fulfil this assumption by introducing more accurate techniques to model both the power consumption and the uncertainty of estimates.

There are many different features which have been investigated that affect power and energy consumption. Mei et al. [9] show that acceleration and deceleration substantially increase power requirements. The impact of specific terrain features is analysed in [10] and [11]. Ishigami et al. in [12] consider terrain elevation and the steering and driving manoeuvres necessary to follow a particular path. Energy-efficient path planning based on estimating terrain traversability is investigated in [13].

While we share the aspiration of some of these related works, our approach is set apart by our intentional reliance on publicly available data. This allows us to predict energy consumptions on any route for which OSM has data.

III. TECHNICAL APPROACH

We define our routes as a collection of connected path segments. The path network and the characteristics of the segments are obtained from collaborative projects, such as OpenStreetMaps (OSM) [14]. Each segment in the OSM map is defined by a feature vector $s \in \mathbb{R}^D$, containing information such as the number of intersections with other segments or surface type. A full explanation of chosen features is given in Section IV. Let p_i denote the power consumption while traversing segment s_i . In order to account for variability from latent factors such as traffic or weather conditions, we implement a heteroscedastic Gaussian Process, following [1]. This uses two different GPs to model the power consumption and the variance of the measurements. The goal of the regression is to infer the posterior distribution $p(P^* | s^*, \mathcal{D})$ of using power P^* to traverse an arbitrary segment s^* , given a training set $\mathcal{D} = \{(s_i, p_i)\}_{i=1}^n$ of n power measurements $\mathbf{P} = (p_1, \dots, p_n)^T$ collected traversing training segments $\mathbf{s} = (s_1, \dots, s_n)^T$. An initial GP (GP₁) is used to

estimate the noise-free power consumption. The empirical noise is modelled by a second GP (GP₂). A third GP (GP₃) is calculated which incorporates both the noise free estimates and the empirical noise. A convergence process then takes place which iterates until all empirical noise is accounted for. The approach is outlined in Algorithm 1.

A. Estimating noise-free power consumption

Our algorithm initially fits a standard homoscedastic GP to the training data:

$$P(\mathbf{s}) \sim \text{GP}_1(m_1(\mathbf{s}), k_1(\mathbf{s}, \mathbf{s}')) \quad (1)$$

where $m(\mathbf{s})$ and $k(\mathbf{s}, \mathbf{s}')$ are the mean and covariance functions chosen to reflect prior information about the data and to maximise the likelihood. GP₁ is used to predict the average noise-free power required to traverse $\mathbf{s}^* = (s_1, \dots, s_q)^T$, an arbitrary vector of q segments:

$$\mu_1(\mathbf{s}^*) = M^* + K^* K^{-1}(\mathbf{P} - M), \quad (2)$$

$$\Sigma_1(\mathbf{s}^*) = K^{**} - K^* K^{-1} K^{*T}. \quad (3)$$

In Equations 2 and 3, we have:

$$M \in \mathbb{R}^n, M_i = m_1(\mathbf{s}_i),$$

$$M^* \in \mathbb{R}^q, M_i = m_1(\mathbf{s}_i^*),$$

$$K^* \in \mathbb{R}^{q \times n}, K_{ij}^* = k_1(\mathbf{s}_i^*, \mathbf{s}_j),$$

$$K \in \mathbb{R}^{n \times n}, K_{ij} = k_1(\mathbf{s}_i, \mathbf{s}_j),$$

$$K^{**} \in \mathbb{R}^{q \times q}, K_{ij}^{**} = k_1(\mathbf{s}_i^*, \mathbf{s}_j^*).$$

B. Modelling Empirical Noise

After fitting the GP₁, the residual noise r_i can be found on each segment s_i in the training segment set \mathbf{s} :

$$r_i = \frac{1}{n_i} \sum_{j=1}^{n_i} (P_j - \mu(s_i))^2 \quad (4)$$

where n_i is the number of power observations associated with segment s_i . GP₂ is fitted to the logarithm of the residual noise, $z_i = \log(r_i)$,

$$z(\mathbf{s}) \sim \text{GP}_2(m_2(\mathbf{s}), k_2(\mathbf{s}, \mathbf{s}')). \quad (5)$$

The mean μ_2 and covariance function Σ_2 for GP₂ are computed in a similar manner to GP₁ using standard GP equations (Equations 2 and 3). The most-likely heteroscedastic framework uses the mode of GP₂ to estimate the noise for any given segment:

$$r(\mathbf{s}) = \exp(\mu_2(\mathbf{s})). \quad (6)$$

As all distributions are Gaussian, the mode of this distribution is the mean. The heteroscedastic model incorporates uncertainty about the measurement variance into a third GP model GP₃ with the posterior distribution:

$$p(P^* | s^*, \mathcal{D}) \sim \mathcal{N}(\mu^*, \Sigma^*) \quad (7)$$

where

$$\mu^* = M + K^*(K + R)^{-1}(\mathbf{P} - M), \quad (8)$$

$$\Sigma^* = K^{**} + R^* - K^*(K + R)^{-1}K^{*T}. \quad (9)$$

The notation defined for Equations 2 and 3 is augmented by:

$$R = \text{diag}(\mathbf{r}) \text{ with } \mathbf{r} = (r(s_1), \dots, r(s_n))^T,$$

$$R^* = \text{diag}(\mathbf{r}^*) \text{ with } \mathbf{r}^* = (r(s_1^*), \dots, r(s_n^*))^T.$$

The algorithm converges by setting $\text{GP}_1 \leftarrow \text{GP}_3$ and iteratively updating the empirical noise. The posterior distribution in Equation 7 can be given to the planner or end user to make informed decisions based on power consumption estimates in different parts of the map.

In this framework, it is possible to use several covariance and mean functions. In our implementation, we use a Matern 5/2 kernel and linear mean function for both GP_1 and GP_2 .

Algorithm 1 Power Consumption Prediction

Input:

$$\mathcal{D} = \{(s_i, p_i)\}_{i=1}^n \quad \triangleright \text{Data Input}$$

$$\mathbf{s}^* \quad \triangleright \text{Segments to calculate Posterior for}$$

Output: μ^*, Σ^* \triangleright Posterior of required power consumption at segments \mathbf{s}^*

Calculate GP₁

- 1: $P(\mathbf{s}) \sim \text{GP}_1(m_1(\mathbf{s}), k_1(\mathbf{s}, \mathbf{s}'))$
- 2: $\mu_1(\mathbf{s}^*) = M^* + K^*K^{-1}(\mathbf{P} - M)$
- 3: $\Sigma_1(\mathbf{s}^*) = K^{**} - K^*K^{-1}K^{*T}$
- 4: **repeat**

Calculate GP₂

- 5: $r_i = \frac{1}{n_i} \sum_{j=1}^{n_i} (P_j - \mu_1(s_i))^2$
- 6: $\mathbf{z} = \log(\mathbf{r})$
- 7: $z(\mathbf{s}) \sim \text{GP}_2(m_2(\mathbf{s}), k_2(\mathbf{s}, \mathbf{s}'))$

Calculate GP₃

- 8: $\mu^*(\mathbf{s}^*) = M^* + K^*(K + R)^{-1}(\mathbf{P} - M)$
- 9: $\Sigma^*(\mathbf{s}^*) = K^{**} + R^* - K^*(K + R)^{-1}K^{*T}$

10: **if** unconverged **then**

11: $\text{GP}_1 \leftarrow \text{GP}_3 \quad \triangleright$ So that $\mu_1(\mathbf{s}^*) = \mu^*(\mathbf{s}^*)$ and $\Sigma_1(\mathbf{s}^*) = \Sigma^*(\mathbf{s}^*)$

12: **end if**

13: **until** convergence

IV. EXPERIMENTAL SETUP

To evaluate the accuracy of the power and energy consumption predictions, we collected data while manually driving a mobile robot in the centre of Oxford. The robot was equipped with current and voltage sensors and a Global Positioning System (GPS) receiver.

A. Platform

The robotic platform used in these experiments is a custom made platform developed within the EC-funded project EUROPA2.0 [15]. The robot is designed to navigate autonomously in urban environments and provide assistance to pedestrians.

B. Map Features

Path segment information is extracted from OpenStreetMaps and GoogleMaps. The complete list of features

used in the predictive model is shown in Table II. Section II describes the motivation for choosing these features. We include the number of surrounding amenities as a proxy for pedestrian traffic. This congestion would increase power consumption as the vehicle stops and starts to avoid people. By the same logic, we consider the presence of traffic lights and pedestrian crossing as they might also cause the robot to stop and start again to cross the street.

Our dataset includes five different routes (shown in Figure 2 and summarised in Table I).

Route	Description	Length (km)
Blue	Passes through a pedestrian area with a high number of amenities	2.59
Orange	Quieter pedestrian route, low traffic	1.62
Purple	Long section along residential streets & gravel paths in park	2.30
Red	Open gravel paths through a park	2.89
Green	Residential streets, short park section and multiple crossings	1.51

TABLE I

SUMMARY OF TRAINING ROUTES

C. Data Collection

Our data collection process took place over two months and recorded over 30km of driving data in a range of different conditions. In total there were 1720 average power measurements on 820 unique segments, representing a broad spectrum of the type of paths in an urban environment would be expected to cover.

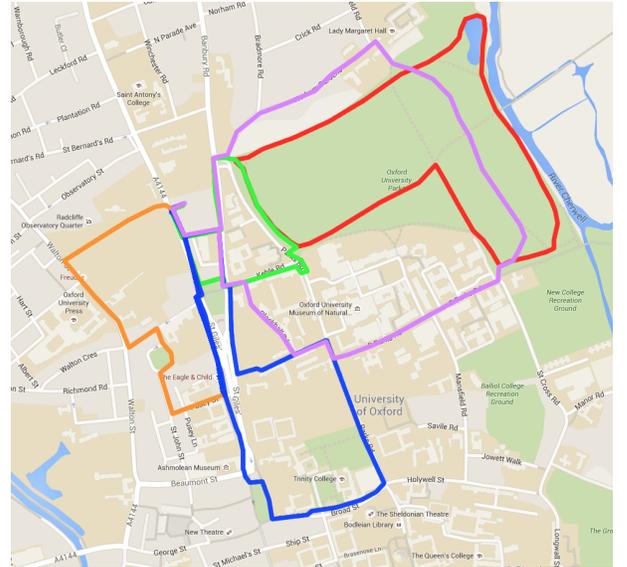


Fig. 2. The five routes traversed by the robot and considered in the experiments.

D. Sensor Data Post-Processing

Current and voltage measurements were collected using an Arduino Micro at a sampling rate of 4KHz and later downsampled to 40Hz. Recorded sensor data are notoriously noisy, which arises from a number of different sources. These

FEATURE NAME	DESCRIPTION	SOURCE
SURFACE (SRF)	terrain type (asphalt, gravel, ..)	OSM
PATH TYPE (PT)	path type (sidewalk, pedestrian area, ..)	OSM
INTERSECTION (INTN)	number of intersections along the path segment	OSM
CROSSING FLAG (CF)	presence of traffic lights or crossing facilities	OSM
AMENITIES (AMN)	surrounding amenities within 100m (museums, pubs, shops, ..)	OSM
ELEVATION (EL)	change of elevation along the path segment	Google Maps

TABLE II

DESCRIPTION OF THE FEATURES USED TO CHARACTERISE THE PATH SEGMENTS AND BUILD THE PREDICTIVE MODEL.

errors can be either systematic (a change in the environmental conditions, such as temperature) or due to random effects (see [16]). Several techniques have been proposed in the literature to post-process sensor data and obtain reliable measurements (see [17], [18]). Following [18], we apply a median filter followed by a mean filter to the current sensor readings to remove outliers and random noise from the data. For both the median and mean filters, we empirically chose a window of five samples. The power consumption p_i for each segment s_i is given by the average power recorded on that segment after processing. Noisy GPS measurements are matched to map segments using the Hidden Markov Model (HMM) based approach proposed in [19]. The power measurements are associated with the respective segment using the timestamps of the GPS readings and sensors. The power consumption of the entire route is obtained by averaging the power consumption of each segment along the route.

V. RESULTS

To evaluate our system, we perform three different experiments. We first analyse the performance of the heteroscedastic regression which predicts power consumption on unseen paths. Later, we introduce a simple benchmark model and compare the predictions provided by such an approach with those provided by our model. This highlights the pertinence of map features for predicting power consumption. Lastly, we evaluate the relevance of those individual features within our model.

A. Energy Consumption Predictions via Heteroscedastic Regression

The purpose of this experiment is to evaluate the accuracy of the power and energy consumption predictions provided by the heteroscedastic regression model on routes that are not observed in training. To do so we perform leave-one-out cross-validation. We train our model with data from every traversal of every route, excluding the route under test, and we compare our predictions against the measured power from all traversals of this route. We measure our performance using the relative error between the predictions and measurements. The results are shown in Table III. Our method provides estimates with an average relative error of 3.31% on the average power estimate and an average relative error of 49.5% on the predicted variance.

As well as performing well on each route, the model also obtains accurate power consumption estimates on a segment

Route	P Measured [W]	P Predicted [W]	μ error [%]	σ error [%]
Blue	561.97 \pm 4.62	557.90 \pm 4.32	0.7	6.5
Orange	580.92 \pm 10.09	556.05 \pm 3.52	4.2	65
Purple	541.10 \pm 19.59	549.48 \pm 2.85	1.5	85.4
Red	565.46 \pm 2.21	539.64 \pm 2.09	4.6	5.4
Green	531.42 \pm 23.18	560.37 \pm 3.44	5.4	85.1
Average	-	-	3.31	49.5

TABLE III

COMPARISON OF MEASURED AND PREDICTED POWER CONSUMPTION FOR DIFFERENT ROUTES, USING THE HETEROSCEDASTIC GP APPROACH.

by segment basis. Figure 3 illustrates the behaviour of our regression model on the five routes individually, showing the power consumption predictions for each segment in each route. We can see that the high variance in the data is well captured by the model and the majority of the power measurements lie within two standard deviations. In some cases, there is bias in the mean prediction that we observe which is due to the evaluation framework chosen. By performing leave-one-out cross-validation, the training set is not always representative of the route travelled. For instance, the red route consists of many gravel segments which are removed by leave one out cross validation, leading to the under-estimate of required power.

Accurate energy estimates for each segment can be calculated by combining our model with a speed profile. We implement a basic speed profile by assuming that the robot always travels at a constant speed, v_{avg} , corresponding to the average speed recorded on the training routes. The energy required to traverse segment i of length l_i is

$$\hat{E}_i = P_i \times \frac{l_i}{v_{\text{avg}}} \quad (10)$$

The results are shown in Table IV which compares the energy consumption measurements E and the energy consumption estimates \hat{E} , calculated by summing all segments in the route. The relative error on the average energy consumption is 6.6%. This suggests that, even with a naive speed profile, it is possible to obtain accurate estimates of the average energy consumption. More advanced speed profiles can only improve this accuracy.

B. The Benchmark

In order to demonstrate that the map features contain information relevant for predicting power consumption, we compare the proposed framework with a benchmark approach. The benchmark approach assumes there are no

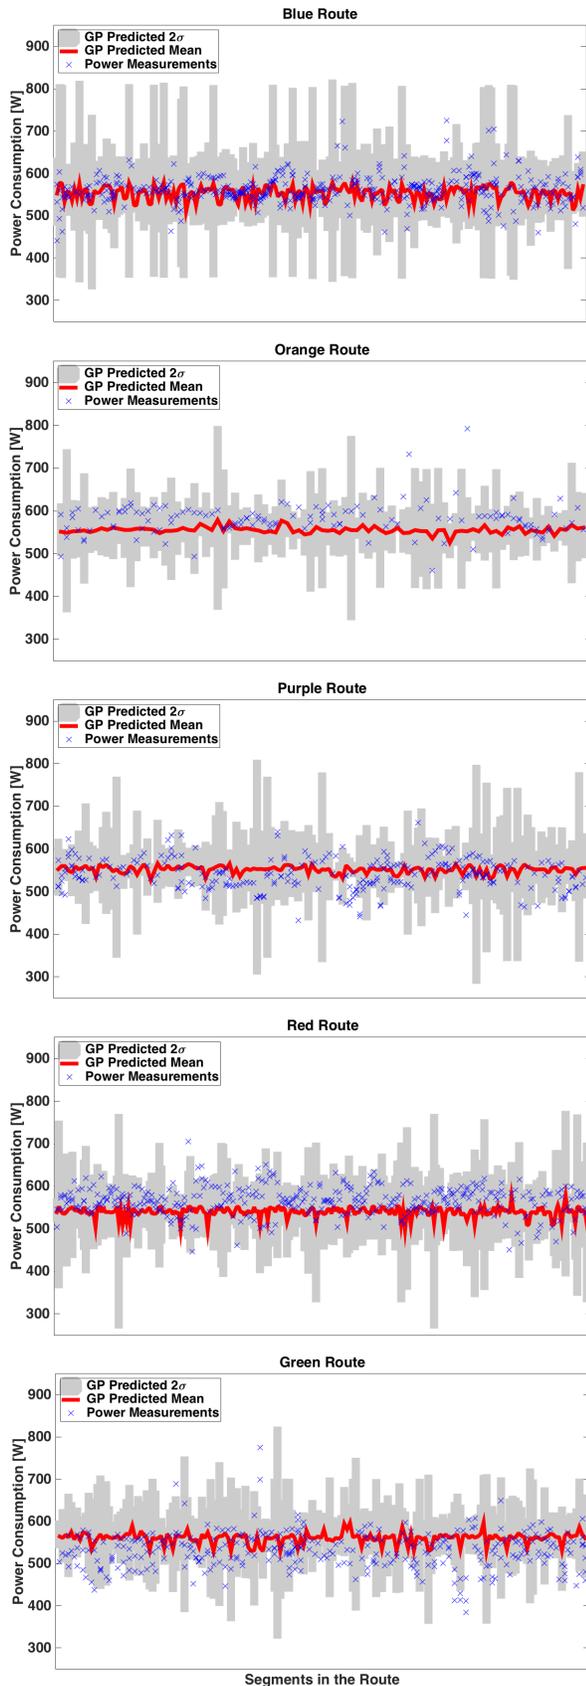


Fig. 3. Power consumption predictions against actual power measurements for each segment of each route. The continuous red line indicates the mean predictions provided by the model, the grey shade indicates the predicted 2 standard deviations. Actual power measurements are indicated by blue crosses.

Route	E [Wh]	\hat{E} [Wh]	μ error [%]
Blue	394.07 ± 7.61	393.93 ± 2.82	0.03
Orange	294.94 ± 25.35	366.4 ± 2.17	24.2
Purple	354.99 ± 54.36	360.90 ± 1.74	1.66
Red	628.78 ± 125.93	627.51 ± 2.07	0.2
Green	260.29 ± 29.97	278.72 ± 1.70	7
Average	-	-	6.64

TABLE IV

COMPARISON OF MEASURED AND PREDICTED ENERGY CONSUMPTION FOR DIFFERENT ROUTES, USING THE HETEROSEDASTIC GP APPROACH.

significant differences amongst the power requirements of different path segments and estimates the power consumption of unknown routes without taking the segments' characteristics into consideration. The benchmark approach uses the mean and variance of the traversed routes as the mean and variance prediction for the unseen route, making the assumption that all segments are the same. We evaluate the behaviour of this model in a leave-one-out cross-validation framework. The obtained results are shown in Table V. When compared to Table III, Table V demonstrates the difference in performance between the benchmark approach and the proposed framework. The relative error on the estimates

Route	P Measured [W]	P Predicted [W]	μ error [%]	σ error [%]
Blue	561.97 ± 4.62	550.14 ± 24.82	2.10	436.7
Orange	580.92 ± 10.09	548.11 ± 20.88	5.6	106.9
Purple	541.10 ± 19.59	555.87 ± 22.80	2.7	16.4
Red	565.46 ± 2.21	549.32 ± 24.38	2.9	1003.2
Green	531.42 ± 23.18	560.13 ± 17.76	5.4	23.38
Average	-	-	3.75	317

TABLE V

COMPARISON OF MEASURED AND PREDICTED POWER CONSUMPTION FOR DIFFERENT ROUTES, USING THE BENCHMARK.

of the mean power using the benchmark method is 3.75%, while the average relative error on the variance prediction is 317%. Our proposed framework shows an improvement of 11% in the accuracy of the prediction of the mean, and an improvement of 84.3% in the accuracy of the prediction of the variance. We attribute this improvement to the fact that the benchmark approach does not take into account the specific characteristics of the segments, but considers all routes and segments equivalent. By incorporating knowledge about the features of each path segment, our proposed methodology is able to model the power consumption and its variance with higher accuracy.

C. Map Feature Ranking

Having shown the information content in the map features, the last experiment aims to evaluate the predictive power of each feature in the GP model. In a leave-one-out cross validation framework, we use the heteroscedastic regression approach to estimate the power consumption using only a subset of the features. We compare the predictions made with all the available features (as a baseline) with predictions made after removing individual features to judge their

predictive power. The decrease in accuracy is expressed as the relative error between the two predictions. Figure 4 shows the result of this analysis. Excluding any of the features has a negative impact on the accuracy of the power consumption predictions; the most discriminant features in our framework are the number of surrounding amenities and the path type. This corroborates our experiences whilst testing, two major reasons for stopping and starting were to avoid pedestrians and to manoeuvre through narrow, awkward streets.

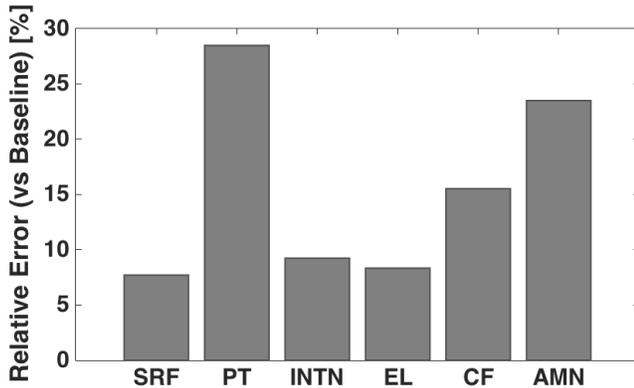


Fig. 4. Comparison between the relative error on the prediction of the variance, when different subsets of features are available. On the y-axis, the relative error is computed with respect to the one given when all features are considered (cfr. Table III). The labels on the x-axis indicate the removed feature (cfr. Table II for full explanation).

VI. CONCLUSIONS

This work addresses the problem of estimating energy consumption of a mobile robot moving in an urban environment. Previous studies investigate how energy consumption is influenced by the kinematics of the robot, its processing and perception. However, using high-level features extracted from publicly available maps, we provide a probabilistic model that accurately predicts energy consumption. As our feature selection is open source, we can extrapolate our predictions to routes that have not been seen previously. Furthermore, as the model does not require information about the internal configuration of the robot, it is applicable to a wide range of platforms. The proposed framework is implemented with Heteroscedastic Gaussian Processes to better represent the variance in the data, which is difficult to model directly. The distinct advantage of our model is our ability to predict measurement variance. Our estimates outperform benchmarks by 84.3%. Such accurate estimates are particularly useful for any planning system that needs to make confident decisions while dealing with limited robot power supplies.

Further investigation could focus on enriching the feature set by extracting information from other sources which might be better predictors of robot efficiency. For example, using Google Streetview data or disabled access data could reveal the state and upkeep of pavements that the robot traverses.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support from the UK EPSRC CDT in AIMS (grant EP/L015987/2) and the

EU project EUROPA2 (FP7-610603). Thanks to Michael Osborne for advice on Gaussian Process regression.

REFERENCES

- [1] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard, "Most likely heteroscedastic gaussian process regression," in *Proceedings of the 24th international conference on Machine learning*. ACM, 2007, pp. 393–400.
- [2] Y. Mei, Y.-H. Lu, Y. C. Hu, and C. G. Lee, "Energy-efficient motion planning for mobile robots," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 5. IEEE, 2004, pp. 4344–4349.
- [3] —, "Deployment of mobile robots with energy and timing constraints," *Robotics, IEEE Transactions on*, vol. 22, no. 3, pp. 507–522, 2006.
- [4] C. H. Kim and B. K. Kim, "Energy-saving 3-step velocity control algorithm for battery-powered wheeled mobile robots," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*. IEEE, 2005, pp. 2375–2380.
- [5] Y. Mei, Y.-H. Lu, Y. C. Hu, and C. G. Lee, "A case study of mobile robot's energy consumption and conservation techniques," in *Advanced Robotics, 2005. ICAR'05. Proceedings., 12th International Conference on*. IEEE, 2005, pp. 492–497.
- [6] S. Fallah, B. Yue, O. Vahid-Araghi, and A. Khajepour, "Energy management of planetary rovers using a fast feature-based path planning and hardware-in-the-loop experiments," *Vehicular Technology, IEEE Transactions on*, vol. 62, no. 6, pp. 2389–2401, 2013.
- [7] J. Brateman, C. Xian, and Y.-H. Lu, "Energy-efficient scheduling for autonomous mobile robots," in *Very Large Scale Integration, 2006 IFIP International Conference on*. IEEE, 2006, pp. 361–366.
- [8] P. Ondruska, C. Gurau, L. Marchegiani, C. H. Tong, and I. Posner, "Scheduled perception for energy-efficient path following," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, May 2015.
- [9] Y. Mei, Y.-H. Lu, C. G. Lee, and Y. C. Hu, "Energy-efficient mobile robot exploration," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*. IEEE, 2006, pp. 505–511.
- [10] Z. Sun and J. Reif, "On energy-minimizing paths on terrains for a mobile robot," in *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, vol. 3. IEEE, 2003, pp. 3782–3788.
- [11] J. A. Broderick, D. M. Tilbury, and E. M. Atkins, "Characterizing energy usage of a commercially available ground robot: Method and results," *Journal of Field Robotics*, vol. 31, no. 3, pp. 441–454, 2014.
- [12] G. Ishigami, K. Nagatani, and K. Yoshida, "Path planning and evaluation for planetary rovers based on dynamic mobility index," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. IEEE, 2011, pp. 601–606.
- [13] S. Martin and P. Corke, "Long-term exploration & tours for energy constrained robots with online proprioceptive traversability estimation," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 5778–5785.
- [14] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *Pervasive Computing, IEEE*, vol. 7, no. 4, pp. 12–18, 2008.
- [15] "The European Pedestrian Assistant 2.0," <http://europa2.informatik.uni-freiburg.de/>, 2013.
- [16] E. Elnahrawy and B. Nath, "Cleaning and querying noisy sensors," in *Proceedings of the 2nd ACM international conference on Wireless sensor networks and applications*. ACM, 2003, pp. 78–87.
- [17] S. Basu and M. Meckesheimer, "Automatic outlier detection for time series: an application to sensor data," *Knowledge and Information Systems*, vol. 11, no. 2, pp. 137–154, 2007.
- [18] M. Weiss, A. Helfenstein, F. Mattern, and T. Staake, "Leveraging smart meter data to recognize home appliances," in *Pervasive Computing and Communications (PerCom), 2012 IEEE International Conference on*. IEEE, 2012, pp. 190–197.
- [19] 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*. ACM, 2009, pp. 336–343.