Real-time energy-efficient path planning for unmanned ground vehicles using mission prior knowledge

Amir Sadrpour* and Jionghua Jin

Department of Industrial and Operations Engineering,
University of Michigan,
1205 Beal Avenue, Ann Arbor, MI 48109-2117, USA
Email: sadrpour@umich.edu
Email: jhjin@umich.edu
*Corresponding author

A. Galip Ulsoy

Department of Mechanical Engineering,
University of Michigan,
2350 Hayward Street, Ann Arbor, MI 48109-2125, USA
Email: ulsoy@umich.edu

Abstract: Unmanned Ground Vehicle (UGV) missions include situations where a UGV has to choose between alternative paths, and are often limited by the available on-board energy. Thus, we propose a dynamic energy-efficient path planning algorithm that integrates mission prior knowledge with real-time sensory information to identify the most energy-efficient path for mission completion. Our proposed approach predicts and updates the distribution of the energy requirement for alternative paths using recursive Bayesian estimation through two stages: (a) exploration – road segments can be explored to reduce their energy prediction uncertainty; (b) exploitation – the most reliable path is selected using the collected information in the exploration stage and then traversed. Our simulation results show that the proposed approach outperforms offline methods, as well as a method that relies on exploitation only to identify the most energy-efficient path.

Keywords: path planning; unmanned ground vehicles; mission prior knowledge; Bayesian estimation; exploration and exploitation.

Reference to this paper should be made as follows: Sadrpour, A., Jin, J. and Ulsoy, A.G. (2014) 'Real-time energy-efficient path planning for unmanned ground vehicles using mission prior knowledge', *Int. J. Vehicle Autonomous Systems*, Vol. 12, No. 3, pp.221–246.

Biographical notes: Amir Sadrpour is a fifth-year PhD student in the Department of Industrial and Operations Engineering at the University of Michigan. He received his MS Degree from the University of Michigan in 2010, and his BS Degree from the Ohio State University in 2008. He joined the Ground Robotic Reliability Center (GRRC) at the University of Michigan in 2010, his research focuses on Unmanned Ground Vehicles (UGVs) reliability modelling, prediction and acceptance testing.

Jionghua (Judy) Jin is a Professor of Industrial and Operations Engineering at the University of Michigan. She received her PhD from the University of Michigan in 1999. Her research focuses on data fusion for complex system improvement, which includes variation reduction, condition monitoring and fault diagnosis, process control, knowledge discovery and decision-making. She has received a number of awards including the NSF CAREER Award in 2002, the PECASE Award in 2004, and Best Paper Awards during 2005–2012. She is a fellow of ASME and a member of ASQC, IEEE, IIE, INFORMS, and SME.

A. Galip Ulsoy is the C.D. Mote Jr. Distinguished University Professor of Mechanical Engineering and the William Clay Ford Professor of Manufacturing at the University of Michigan. He received the PhD from University of California at Berkeley in 1979, the MS Degree from Cornell University in 1975, and the BS Degree from Swarthmore College in 1973. His research interests are in the dynamics and control of mechanical systems. He has received numerous awards, including 2008 Rufus T. Oldenburger Medal from ASME. He is a member of the National Academy of Engineering and is a Fellow of ASME, SME and IFAC.

This paper is a revised and expanded version of a paper entitled 'Real-time energy-efficient path planning for unmanned ground vehicles using mission prior knowledge' presented at the 'Dynamic Systems and Control Conference', Stanford, CA, October 2013.

1 Introduction

1.1 Problem statement

Unmanned Ground Vehicles (UGVs) are entering the economic mainstream and are now being used extensively in military and commercial applications (Tilbury and Ulsoy, 2011). Even with a rapid increase in the number of UGVs, several studies show average failure rates (every 6–20 h) much worse than the 96 h benchmark established by the Department of Defense (Kramer and Murphy, 2006). Despite their sub-optimal reliability (Carlson and Murphy, 2003; Carlson and Murphy, 2005; Carlson et al., 2004; Stancliff and Dolan, 2005; Sadrpour et al., 2011), UGVs are expected to be a safe and cost-efficient option for space missions, rescue operations and military applications.

One of the key factors that limit the utility of small tele-operated battery-powered UGVs in surveillance missions is the available on-board energy. The vehicle locomotion is the main source of energy consumption for most UGVs (Sadrpour et al., 2013a). Typical mission duration is currently on the order of 1–2 h, while it is often desirable to carry out much longer missions (e.g. 8–10 h) between lengthy recharging stops. A typical surveillance mission consists of various tasks and several alternative paths. Due to limited energy storage capacity, it is essential to predict the energy requirement of alternative paths to help the operator with path planning. The goal of surveillance missions studied here is to start from a known location on a map and reach a destination

point using one of the available alternative paths. The objective of this research is to identify the path with the highest probability of successfully completing the mission using the information available at any given time. One failure mode of interest is the unanticipated depletion of the UGV's stored energy, which results in failures to reach the destination point. The shortest path is not always the most energy-efficient since in addition to length, other factors such as road roughness and grade and driving style affect the energy consumption. Additionally, the recommended criterion of the lowest failure probability, instead of the minimum expected energy consumption, considers the prediction uncertainty as well as the expected path energy requirement in decision-making.

The rest of the paper is organised as follows: Section 1.2 provides an overview of the literature in the area of path planning with energy consideration, which helps justify the contribution of our research. Section 2 is an overview of the proposed methodology. Section 3 describes in detail the proposed Bayesian algorithm for the most energy-efficient path planning. Section 4 presents a comparative simulated case study to illustrate the advantages of the approach. Concluding remarks are given in Section 5.

1.2 Literature review and contributions

In this paper, a surveillance mission is represented by a network where arcs symbolise road segments and nodes represent intersections of road segments. The cost of each arc is the energy required to traverse the arc. This energy requirement is affected by variable random factors such as road surface conditions and grades with unknown probabilistic distributions. Mei et al. (2005) measured the power consumption of different components of UGVs and presented strategies for saving of energy that take advantage of UGV idle time, speed, etc. However, their deterministic energy models cannot consider random road grade variations, and no case study is presented on applying such strategies for real-time path planning.

Our problem falls within the general class of Shortest Path Problems (SPP). Deterministic SPPs (Denardo, 2003) as well as stochastic shortest path problems (SSPP) with known cost distributions (Powell, 2011; Fan et al., 2005) have been extensively studied. A stochastic most reliable path problem with normally and correlated random costs were investigated by Seshadri and Srinivasan (2010); however, the distributions of costs were assumed to be precisely known prior to the mission. Our paper relaxes the assumption of known path cost distributions, and further considers the uncertainty of path costs in the planning stage.

When the distributions of paths' costs are not known, adaptive learning via exploration becomes a viable approach in decision-making. Exploration is a process by which an arc cost distribution is estimated more precisely by collecting actual operating data for a short period from the arc. Ryzhov and Powell (2011) introduced an exploration policy based on the Knowledge Gradient (KG) in a stochastic SPP with unknown cost distributions, in which exploration could be performed on any arc in the network at any given time. In contrast, UGV can only collect measurements from the sequential road segments that it traverses. Also, our exploration cost increases when additional sampling information is needed.

Several papers investigate variations of the Travelling Salesman Problem (TSP) (Applegate et al., 2011) for UGV path planning with energy consideration. In the work of Wei et al. (2012), a path planning problem was discussed for mobile robots with the objective of minimising the energy requirement using docking stations with deterministic arc costs. A TSP for mobile robots was considered in the study of Sipahioglu et al. (2008) with dynamically changing paths using a deterministic cost model. Moreover, energy-based path planning for cabled robots was studied by Borgstrom et al. (2008). Their goal was to maximise the accumulated rewards by visiting a sequence of nodes in a network. A similar concept is utilised in our proposed approach during the exploration stage. We prefer road segments whose exploration yields the maximum reduction of the prediction uncertainty by considering a stochastic cost model.

Another class of energy-efficient path planning for small UGVs deals with a coverage task problem. In a coverage task, the UGV is required to move through an area and travel within a certain distance of predefined way-points. Broderick et al. (2012) investigated an energy-efficient coverage task using optimal control. Unlike a shortest path formulation, the UGV must visit every point on the map. Coverage tasks were also studied in the work of Mei et al. (2004) with a focus on minimising the energy for locomotion by optimally tuning the vehicle velocity and trajectory. However, their models were deterministic and did not consider the impact of terrain variations on power consumption.

Reinforcement learning (RL) is a class of online learning approaches, where an agent interacts with a stochastic and dynamic environment and learns a policy to maximise a measure of its long-term reward (Sutton and Barto, 1998; Dearden et al., 1999; Dearden et al., 1998). Many RL approaches deal with the trade-off between exploration and exploitation. There are three major differences between those traditional RL frameworks and the proposed learning scheme to be discussed in this paper: (a) in our problem, the risk associated with exploration grows with additional measurements due to limited on-board energy; (b) since there is no inherent exploration risk in RL when dealing with Markov decision processes, most exploration strategies have an oscillatory behaviour in which alternatives or states are visited in an alternating fashion. With a UGV, due to physical constraints, such exploratory strategies are not energy-efficient; (c) the reward function in our paper uses the criterion of lowest failure probability (highest reliability) that considers both the expected energy requirement of road segments and their covariance, resulting in a reward structure that is not independent of past or future states of the vehicle. Seshadri and Srinivasan (2010) showed, through a counterexample, how the inclusion of covariance in the structure of the reward function results in inapplicability of traditional shortest path algorithms for finding the most reliable path.

The objective of this paper is to present a novel path planning problem for UGVs, under a network of alternative paths with the following characteristics: (a) the arc cost distributions are not precisely known *a priori*; (b) the arc costs may be correlated; (c) the distributions of arc costs can be updated online based on real-time measurements; and (d) UGVs can only collect measurements from the road segment that they traverse. To identify an energy-efficient path in the network, we propose a heuristic approach that integrates mission prior knowledge and real-time measurements for adaptively predicting the energy requirement distributions of alternative paths. The proposed method is described in the next section.

2 Methodology overview

Figure 1a illustrates the framework of the proposed approach. Let us assume the vehicle has reached an intersection (node) in a network (see Figure 1b), from which alternative road segments emanate. A vehicle longitudinal dynamic model and mission prior knowledge are used for estimating the initial distribution (e.g. mean and variance) of energy requirement of alternative paths (Sadrpour et al., 2012). The first step uses the initial distributions to remove paths that are very unlikely to be the most energy-efficient from consideration through a process termed *pruning*. For instance, Figure 2a depicts the initial distributions of three alternative paths of the network in Figure 1b, i.e. $Q=\{q_1, q_2, q_3\}$, where the set q_i contains the indexes (e.g. {*a*, *b*, *c*, …}) of road segments of path *i*, based on mission prior knowledge assuming that the vehicle is at node 1. The figure shows that the distribution of the third path exceeds and does little overlap with the first two paths. If this situation arises, in favour of $\{q_1, q_2\}$, path q_3 can be pruned without a need for exploration. If initial pruning results in only one unpruned path, no further action is needed and UGV can use the remaining path to reach the destination. However, some paths may still remain overlapping after pruning due to large energy prediction uncertainty. In this case, *exploration* of the remaining paths may become necessary, in which some of the available energy is used to explore the remaining alternatives (i.e. traverse and backtrack if necessary) to reduce their prediction uncertainty and bias.

The exploration step includes two sub-steps. First is to evaluate the exploration feasibility, i.e. to determine if exploration of alternative road segments emanating from the current node is feasible considering the predicted energy of the paths and the remaining stored energy in the UGV. If exploration is not feasible, the most fuel efficient path is selected based on the available information. The criterion for ranking and selection of energy-efficient paths is defined by the ratio $z=\left(\frac{\textit{failure threshold}-\textit{predicted energy expectation}}{\textit{prediction standard deviation}}\right)$, which is termed the *z*-score and considers both the (expected) predicted path energy requirement as well as the prediction uncertainty. If exploration is feasible, in sub-step 2, the number of exploration measurements from each road segment, i.e. the exploration budget assignment, is determined using an energy-efficient strategy based on the reduction of the energy prediction uncertainty. By collecting the measurements, the energy distributions are updated. For example, Figure 2b shows the updated distribution of the remaining paths, i.e. $\{q_1, q_2\}$, after exploring road segments {*a*, *b*}. Since the distributions of $\{q_1, q_2\}$ no longer overlap, using the pruning criterion, path q_2 is eliminated in favour of q_1. The exploration may not always lead to one remaining unpruned path. In either case, the road segment of the most energy-efficient path based on the updated energy distributions is selected to be *exploited*.

The exploitation step includes traversing a road segment until the vehicle reaches the next intersection in the network. During exploitation, pruning still continues in real-time. The exploration and exploitation steps are repeated whenever the vehicle reaches a node with multiple alternative paths until it arrives at the destination node. For instance, based

on Figure 2b, path $q_1 = \{a, c\}$ is selected to be exploited since it has the highest efficiency (highest z-score) and because other alternatives paths have been pruned. Section 3 will describe the above steps in detail.

Figure 1 (a) Methodology overview. (b) Example of a simple network with 3 nodes, 4 road segments, start node 1, and end node 3. The alternative paths are $Q=\{q_1=\{a,c\},q_2=\{b\},q_3=\{a,d\}\}$

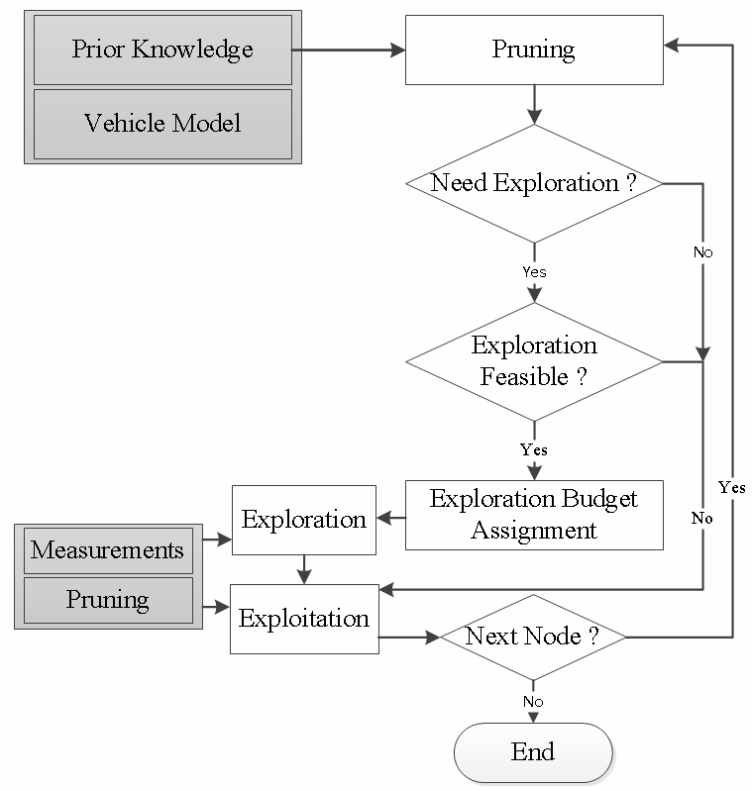

(a)

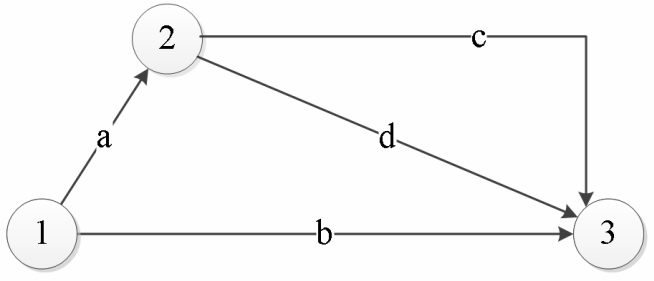

(b)

Figure 2 (a) The initial energy requirement distribution of a hypothetical mission with three alternative paths (see Figure 1b) based on mission prior knowledge. Path q_3 energy requirement is clearly higher than $\{q_1, q_2\}$. However, it is not clear which of the remaining paths is more energy-efficient. (b) The updated distributions of energy requirement of $\{q_1, q_2\}$ after exploring them. Q_2 can be pruned in favour of q_1 which is estimated to have the highest probability of success (see online version for colours)

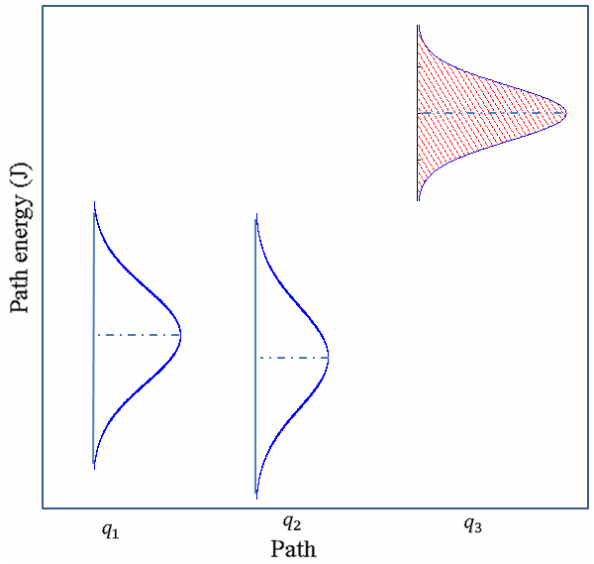

(a)

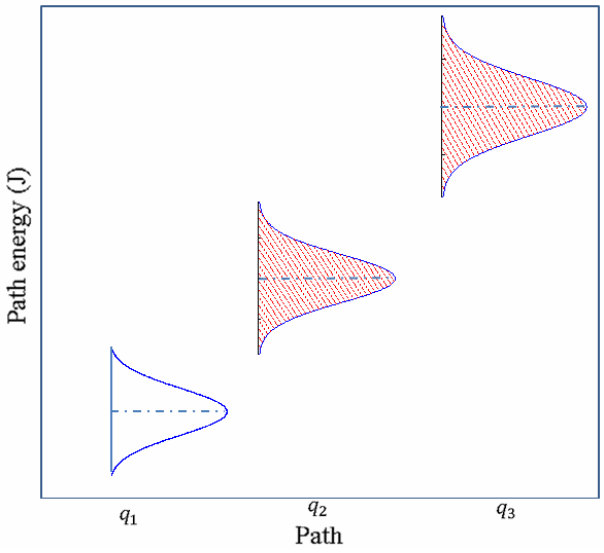

(b)

3 Methodology

3.1 Vehicle model

A linearised vehicle longitudinal dynamics model, as typically used for power consumption studies in automobiles, is also utilised here (Ulsoy et al., 2012). The UGV power consumption is modelled as follows (Sadrpour et al., 2013b):

$$P(t)=\left(W\theta(t)+fW+ma(t)+C_I\right)u(t)+\beta+\varepsilon(t) \tag{1}$$

where $P(t)$ is the power at time t, W is the vehicle weight, $\theta(t)$ is the road grade, f is the road rolling resistance coefficient, m is the vehicle mass, $a(t)$ is the acceleration, C_I is the internal resistance coefficient, $u(t)$ is the velocity, β represents other constant sources of energy depletion, such as electronic sensors on-board the vehicle, and $\varepsilon(t)$ is the model error following NID(0, σ_ε^2). Other time varying factors, which have a smaller relative significance, such as aerodynamic drag, are neglected here due to the low operating speed of small UGVs. Also, according to experimental results presented by Sadrpour et al. (2013a), the vehicle slippage is negligible on steep uphill and downhill roads, and the current vehicle longitudinal dynamics model does not consider slippage.

In practice, the actual instantaneous UGV power consumption can be obtained in real-time by multiplying the measured current and voltage of the battery. The vehicle velocity can also be measured using a wheel velocity encoder. The acceleration can be estimated based on the difference between two consecutive velocity measurements. Generally, the exact values of rolling resistance coefficient, road grade and vehicle internal resistance are difficult to know beforehand; however, some rough knowledge of the vehicle characteristics and road conditions, which can be generally expressed by a prior probability distribution, might be available.

Equation (1) can be rewritten as a linear regression model:

$$y(t)=Cx(t)+\varepsilon(t) \tag{2}$$

where $y(t)=P(t)-ma(t)u(t)-\beta$, $x(t)=u(t)W$ and $C=\theta(t)+f+C_I^{'}$ is the regression model parameter that combines the grade, rolling resistance coefficient and internal frictional losses. For ease of notation and without loss of generality, we define $C_I^{'}=C_I/W$. The proposed vehicle model was validated by experimental studies of Sadrpour et al. (2013a, 2013b).

Parameter C in equation (2) represents the average combined parameter of a road segment. This model (i.e. fixed effect represented by C) does not capture the natural variations in the grade and rolling resistance coefficients within a road segment. Consequently, the prediction variance is underestimated. To overcome this shortcoming, a mixed effect (random slope) model, i.e. $y=\left(C+\mathfrak{C}_\ell\right)x+\varepsilon$, where $\mathfrak{C}_\ell's$ are i.i.d with $\mathfrak{C}_\ell\sim N(0,\sigma^2)$ is used where the estimate of σ^2 can be obtained by experiments such as the ones presented by Sadrpour et al. (2013a) and procedures in Appendix A.1. A road segment is divided into smaller sub-segments $\ell=1,2,\ldots,L$ and n_ℓ measurements are collected from each sub-segment, ℓ. In the mixed effect model, parameter C captures the average combined parameter (slope) and $\mathfrak{C}_\ell$ captures the deviations of each sub-segment slope from the average slope. The mixed effect model provides a more reasonable estimate of the prediction uncertainty compared to the fixed effect model; however,

parameter estimation for the mixed effect model is more complex computationally since the posterior distributions of parameters do not have a closed-form expression (see Appendix A.1). We assume the random effect has a negligible effect on estimation of C, and will use the fixed effect model to update the posterior distribution of C. The random effect model is used to obtain more accurate estimate of the prediction variance, and the model parameter, σ^2, does not need to be updated with every measurement as discussed in Appendix A.2.

3.2 Problem definition

Consider a directed network $G(N, A)$ with nodes $N = \{1, 2,\ldots,n\}$, and edges $A = \{a, b, \ldots\}$, where edges represent road segments. A road segment is defined as a sector of a road that has the same distribution of grade and rolling resistance (e.g. uphill/grass, level/paved, downhill/unpaved). The goal of a UGV operator is to traverse from a starting node to a final node using one of the alternative paths. Our objective is to provide the operator with a path that has the highest probability of reaching the destination without running out of energy.

3.3 Bayesian prediction of energy requirement

The energy requirement of road segment i, i.e. E_i where $i \in A$, can be calculated by integrating the instantaneous power over the time spent on the segment, i.e.

$$E_i = \int_0^{t_e} P(t)dt \approx \sum_{j=1}^{n} P(j)\Delta t \tag{3}$$

where E_i is the total energy requirement of the road segment i, and t_e, n and Δt are the end time, number of measurements, and the sampling interval, respectively. Let us assume that the vehicle has collected k measurements from road segment i. E_i can be estimated by

$$\hat{E}_i(k) = E_i^o(k) + \hat{E}_i^r(k) \tag{4}$$

where $E_i^o(k)$ is the measured actually consumed energy up to time $t = k\Delta t$ and $\hat{E}_i^r(k)$ is the predicted expected energy requirement for the remainder of the segment.

According to equations (3) and (4), to estimate E_i, predictions of power for each road segment in the network are carried out using the vehicle model [(equation (2)]. In the work of Sadrpour et al. (2013a, 2013b), Bayesian recursive estimation was used to estimate and update the unknown parameter C of equation (2) using real-time velocity and power measurements and mission prior knowledge. In equation (2), the prior distribution of C is assumed as follows:

$$C_i^0 = f_i^0 + \theta_i^0 + C_I^{'0} \sim N\left(\mu_{f_i}^0 + \mu_{\theta_i}^0 + \mu_{C_I'}^0, \sigma_{f_i}^{0\,2} + \sigma_{\theta_i}^{0\,2} + \sigma_{C_I'}^{0\;2}\right) \tag{5}$$

where $\mu_{f_i}^0$, $\mu_{\theta_i}^0$ and $\mu_{C_I'}^0$ are the means of the prior distributions of rolling resistance coefficient, average grade and vehicle internal resistance, respectively, for road segment i, and $\sigma_{f_i}^{0\,2}$, $\sigma_{\theta_i}^{0\,2}$ and $\sigma_{C_I'}^{0\;2}$ are the corresponding variances of the prior distributions.

Sadrpour et al. (2013a) experimentally estimated the prior distributions of the parameters for various typical road surfaces using an iRobot Packbot. The chassis is propelled by two tracks on either side of the vehicle as shown in Figure 3. Assuming k measurements of velocity and power have been collected from road segment i, the posterior distribution of C_i in model [(equation (2)] is updated as follows (Congdon, 2003):

$$\hat{C}_i(k|k-1) \sim N\left(\hat{\mu}_i^C(k|k-1), \left[\hat{\sigma}_i^C(k|k-1)\right]^2\right) \tag{6}$$

where

$$\hat{\mu}_i^C(k|k-1) = \left(\left[\hat{\sigma}_i^C(k-1|k-2)\right]^{-2} + \sigma_\varepsilon^{-2} x^2(k)\right)^{-1} \times \left(\left[\hat{\sigma}_i^C(k-1|k-2)\right]^{-2} \hat{\mu}_i^C(k-1|k-2) + \sigma_\varepsilon^{-2} x(k) y(k)\right)$$

$$\left[\hat{\sigma}_i^C(k|k-1)\right]^2 = \left(\left[\hat{\sigma}_i^C(k-1|k-2)\right]^{-2} + \sigma_\varepsilon^{-2} x^2(k)\right)^{-1}$$

where $\hat{\mu}_i^C(k|k-1)$ and $[\hat{\sigma}_i^C(k|k-1)]^2$ represent the k-th update of the mean and variance of C_i, respectively. Note that for real-time prediction, the estimates of C are obtained from the fixed effect model. The CPU time for estimation is negligible, as is the case for most recursive algorithms. The computer holds the estimated posterior mean and variance of C from the last iteration, and updates them using the closed form expression (6). Using the posterior predictive distribution of $y(k)$, we can estimate the distribution of energy requirement of each road segment as follows (Appendix A.2):

$$\hat{E}_i(k) \sim N\left(\hat{\mu}_i^E(k), \left[\hat{\sigma}_i^E(k)\right]^2\right) \tag{7}$$

where

$$\hat{\mu}_i^E(k) = E_i^o(k) + \hat{t}_i(k)\beta_i + \hat{r}_i(k) W \hat{\mu}_i^C(k|k-1)$$

$$\left[\hat{\sigma}_i^E(k)\right]^2 = W^2 \hat{r}_i^2(k)\left[\hat{\sigma}_i^C(k|k-1)\right]^2 + \hat{t}_i(k)\hat{\sigma}_\varepsilon^2 \Delta t + \phi_i \hat{r}_i(k)$$

where $\hat{\mu}_i^E(k)$ and $[\hat{\sigma}_i^E(k)]^2$ represent the k-th update of the mean and variance of $\hat{E}_i$, respectively, $\hat{t}_i(k)$ and $\hat{r}_i(k)$ are the estimated remaining time and remaining distance of road segment i, and ϕ_i is a constant (see Appendix A.2). To obtain an estimate of $r_i(k)$, we assume that real-time localisation is available using Global Positioning System (GPS) or Simultaneous Localisation and Mapping (SLAM) techniques.

We declare two road segments similar if they have the same prior distributions for rolling resistance and grade (i.e. they share the same parameter C). Measurements from one road segment are used to update the energy requirement distribution of all similar road segments in the network (Appendix A.3). The covariance of energy prediction between two similar road segments i and i' is given as follows (Appendix A.4):

$$\hat{\sigma}_{(i,i')}^E(k) = W^2 \hat{r}_i(k) \hat{r}_{i'}(k)\left[\hat{\sigma}_i^C(k|k-1)\right]^2 \tag{8}$$

Figure 3 The iRobot Packbot was used to estimate the prior distributions parameters

The covariance between roads that are not similar is set to zero. Using equations (7) and (8), we construct the joint distribution of energy requirement of road segments. The next step is to enumerate all paths from the start node to the destination node to construct the set $Q = \{q_1, q_2, \ldots\}$. The computational complexity for enumerating all the paths from the start to end node, where all N nodes are connected, is exponential in the number of nodes. However, in real-world applications the computational complexity is expected to be much less since most nodes are not connected and many paths can be pruned against obviously more reliable ones. For example, Figure 6b depicts a network with seven nodes in which the UGV travels across the University of Michigan's north campus from node 1 to node 7. Since a path is composed of one or more road segments, its energy requirement distribution is predicted by the sum of energy requirement distributions of its corresponding road segments as follows:

$$\hat{E}_{q_j}(k) = \sum_{i \in q_j} \hat{E}_i(k) \tag{9}$$

Let us denote $\hat{\mu}^E_{q_j}(k)$ and $\hat{\sigma}^E_{q_j}(k)$ as the estimated mean and standard deviation of the path energy requirement at time k. When energy distributions are not known, $\hat{\mu}^E_{q_j}(k)$ and $\hat{\sigma}^E_{q_j}(k)$ will be updated with each new measurement, which in turn changes the estimated success probability of path q_j. Thus, there may be a situation that while traversing a path with a previously estimated high probability of success, it becomes unreliable. Moreover, some paths that are not initially selected may actually have higher probability of success and measurements may not be collected from them to update their energy distributions. To reduce the impact of the above issues, paths are explored with an energy-efficient strategy.

Let us define $z_{q_j}(k)=\left(\frac{T-\hat{\mu}^E_{q_j}(k)}{\hat{\sigma}^E_{q_j}(k)}\right)$ where T is the available energy prior to the mission. This index is commonly known as the z-score, which is used to reflect the reliability (i.e. probability of success) of a path indicated by $\hat{R}_{q_j}(k)=\Phi\left(z_{q_j}(k)\right)$ where $\Phi(\bullet)$ is the cumulative distribution function of the standard normal distribution. The benefits of using the z-score is that it not only captures the reliability of a path (i.e. a path with a higher reliability has a higher z-score), but it is a more informative measure to identify the minimum energy path with high certainty. To rank paths based on their probability of success, we will use the z-score of the paths.

3.4 Pruning of undesirable paths

Pruning requires calculation of the lower and upper confidence intervals of the energy requirement of each path, i.e. $LCI_{q_j}(k)$, $UCI_{q_j}(k)$, using $\hat{\mu}^E_{q_j}(k)\pm z_\alpha\hat{\sigma}^E_{q_j}(k)$, where z_α, $0\le\alpha\le1$, is the 100(1–α)th percentile of the standard normal distribution. Using a pair-wise comparison of upper and lower confidence intervals if $\{\exists\, q_j, q_{j'}\in Q: LCI_{q_j}(k)> UCI_{q_{j'}}(k)\}$, then path q_j is eliminated from Q. The road segments of the remaining paths may be traversed in the exploration stage.

3.5 Exploration feasibility & exploration budget

Prior to exploring a path, the uncertainty of energy predictions, i.e. $\left\{\hat{\sigma}^E_{q_j}: j\in Q\right\}$, can be very large. As a result, the prediction confidence intervals of energy requirements of many alternatives paths may overlap. The main objective of exploration is to separate some overlapping energy distributions. The separation is a result of reduction in the energy prediction uncertainty and bias from imprecise prior knowledge.

To determine the feasibility of exploration, we need to ensure that energy spent on exploration does not reduce the probability of mission success below a desirable threshold. The exploration feasibility study provides a threshold for the maximum allowable exploration energy expenditure, i.e. an exploration budget.

The exploration budget, i.e. $Ex_{q_j}(k)$, is path dependent. Not all the road segments will be visited during a mission. To obtain the exploration budget, we modify the reliability function $\hat{R}_{q_j}(k)$ of q_j by adding $Ex_{q_j}(k)$ to its expected exploitation energy requirement. For each remaining path after pruning, we determine $Ex_{q_j}(k)$ satisfying the following condition $\Phi\left(\frac{T-\hat{\mu}^E_{q_j}(k)-Ex_{q_j}(k)}{\hat{\sigma}^E_{q_j}(k)}\right)>1-\gamma$. Parameter γ is the probability of

mission failure after conducting both exploration and exploitation if path q_j is ultimately selected for exploitation and is calibrated by the user. Solving the relationship for $Ex_{q_j}(k)$ results:

$$Ex_{q_j}(k)<T-\left(\hat{\mu}^E_{q_j}(k)+z_\gamma\hat{\sigma}^E_{q_j}(k)\right) \tag{10}$$

Since at this stage, any of the remaining paths may potentially be the most energy-efficient, the exploration budget is selected so that even in the worst case scenario the chance of completing the mission is still $1-\gamma$ using $Ex_m(k)=\max\{\min_{q_j}\{Ex_{q_j}(k):\ \forall q_j\in Q\},0\}$. This is the smallest exploration budget found by calculating $Ex_{q_j}(k)$ for all paths in Q.

3.6 Exploration budget assignment

Suppose the vehicle is at a node from which alternative road segments emanate, i.e. an exploratory node. The objective of exploration budget assignment is to determine how many measurements to collect from each road segment emanating from an exploratory node. Let $q^{Ex}_{j,n}$ denote the set of road segments along path q_j that can be potentially explored when UGV is at node n. For instance, based on Figure 1b, $q^{Ex}_{1,1}=\{a,b,c,d\}$ and $q^{Ex}_{1,2}=\{c,d\}$.

Two competing criteria in the exploration budget assignment are: (a) the reduction in $\hat{\sigma}^E_i$ as a result of additional measurements; and (b) the energy consumed for collecting the measurements. Equation (6) provides a closed-form relation for posterior variance of C_i as a function of predictor $x(k)$. $\hat{\sigma}^E_i$ is expected to decrease with more measurements because $\hat{r}_i(k),\left[\hat{\sigma}^C_i(k|k-1)\right]^2,\hat{t}_i(k)$ are all generally expected to decrease with additional measurements. The variance updates depend on the input $x(t)$ (weighted drive cycle), which is not known *a priori*. However, we can still estimate the reduction of $\hat{\sigma}^E_i$ by simulating a drive cycle from a velocity model. In this study, we assumed that the velocity follows a normal distribution $u\sim NID\left(\mu_u,\sigma_u^2\right)$, but other models such as time series can also be used. The expected reduction of $\hat{\sigma}^E_i$ at time k can then be estimated as (see Appendix A.5):

$$\begin{aligned}\left[\breve{\sigma}^E_i(\breve{k}_i)\right]^2=&W^2\left(d_i-\mu_u\breve{k}_i\Delta t\right)^2\left(\left[\hat{\sigma}^C_i(k|k-1)\right]^{-2}+\sigma_\varepsilon^{-2}\breve{k}_iW^2\left(\mu_u^2+\sigma_u^2\right)\right)^{-1}\\&\left(\frac{d_i}{\mu_u}-\breve{k}_i\Delta t\right)\sigma_\varepsilon^2\Delta t+\phi_i\left(d_i-\mu_u\breve{k}_i\Delta t\right)\end{aligned} \tag{11}$$

where $[\breve{\sigma}^E_i(\breve{k}_i)]^2$ represents the simulated variance as a function of $\breve{k}_i$, which is the number of simulated velocity measurements from road segment i, and d_i is the length of road segment i. An estimate for the expected cost of collecting $\breve{k}_i$ measurements from road segment i at time k is obtained by

$$E^{Ex}_i(\breve{k}_i)=2\left(\hat{\mu}^C_i(k|k-1)\mu_uW+\beta\right)\breve{k}_i\Delta t \tag{12}$$

The multiplier 2 in equation (12) is used because we assume exploration is a round trip operation, and the vehicle returns to the exploratory node after exploring a road segment. Figure 4 shows that $f\left(\breve{k}_i\right)=\breve{\sigma}_i^E\left(\breve{k}_i\right)+\xi E_i^{Ex}\left(\breve{k}_i\right)$ is a convex function for a typical road segment where $\xi \geq 0$ represents the relative importance of exploration cost with respect to variance reduction. To assign the exploration budget to $q_{j,n}^{Ex}$, we solve the following optimisation problem for path q_j:

$$\min \sum_{i \in q_{j,n}^{Ex}} \breve{\sigma}_i^E(\breve{k}_i)+\xi E_i^{Ex}(\breve{k}_i) \tag{13}$$

$$\text{subject to: } \sum_{i \in q_{j,n}^{Ex}} E_i^{Ex}\left(\breve{k}_i\right) \leq Ex_m(k), \breve{k}_i \in \mathbb{Z}$$

The decision variable in this optimisation, $\breve{k}_i$, is the number of measurements to collect from road segment *i* during exploration. Carrying out the optimisation for all paths that pass through node *n*, the output is the assignment of the exploration budget to the explorable road segments in the network. Although only the immediate road segments can be explored, future road segments are considered to assure that sufficient energy is available for future exploration. The optimisation is repeated whenever the vehicle reaches an exploratory node.

Figure 4 The convex structure of $f(\breve{k})=E^{Ex}(\breve{k})+\breve{\sigma}^E(\breve{k})$ for a road segment with the following parameters $d=5000$ m, $\Delta t=1$ s, $\hat{\sigma}^C(0)=0.043$, $\hat{\mu}^C(0)=0.3$, $W=400$ N, $\sigma_\varepsilon=7$ W, $\mu_u=1.5$ m/s, $\sigma_u=0.3$ m/s, $\beta=28$ W, $\phi=30{,}000$. The value of $\breve{k}$ corresponding to minimal $f(\breve{k})$ is the ideal number of exploring measurements from the road segment (see online version for colours)

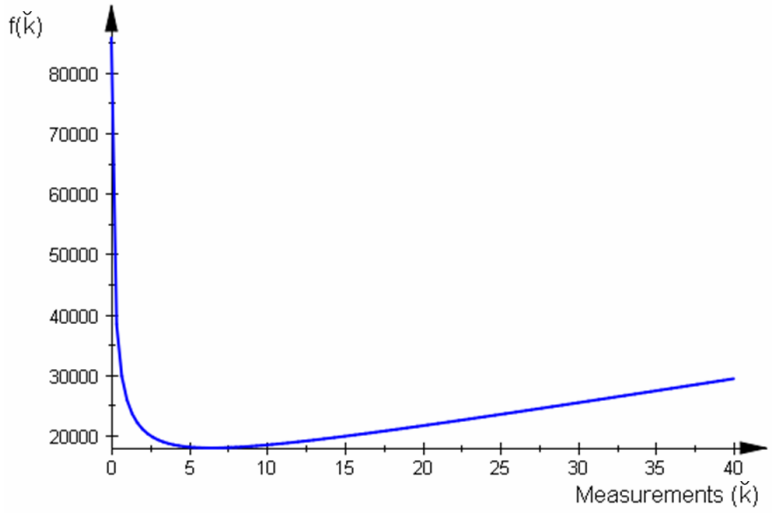

3.7 Exploitation

In the exploitation stage, using the collected information during the mission execution and exploration, the path with the highest z-score, i.e. $\arg\max_{q_j \in Q} z_{q_j}(k)$, is selected to be traversed or exploited. The z-score of all the remaining paths in Q can be calculated as follows:

$$z_{q_j}(k) = \frac{T - \hat{\mu}^E_{q_j}(k) - \sum_{i \in q^{Ex}_{j,n'}} 2\breve{k}_i \left(\hat{\mu}^C_i \left(k|k-1 \right) \hat{x}(k) + \beta \right)}{\hat{\sigma}^E_{q_j}(k)} \tag{14}$$

where $\hat{x}(k)$ is the expected input using a weighted average of past input measurements, and n' is the next node to be visited along path q_j. Equation (14) considers the energy to traverse the path as well as an estimate of the cost of future exploration for road segments along this path, i.e. $\sum_{i \in q^{Ex}_{j,n'}} 2\breve{k}_i (\hat{\mu}^C_i \left(k|k-1 \right) \hat{x}(k) + \beta)$.

4 Simulated case study

In this section, we compare our method with four other approaches for identifying an energy-efficient or reliable path in a network.

1 *Naive approach*: This approach does not consider any model for prediction and relies only on mission qualitative prior knowledge. It uses intuition for comparing the energy requirements of alternative paths. For instance, unpaved roads require more energy per unit distance travelled compared to paved roads.

2 *Minimum expected energy without updating*: This approach assumes the distribution of energy requirement of road segments cannot be updated and are known as *a priori*. Dijkstra's algorithm can be used to identify the optimal path (Denardo, 2003).

3 *Most reliable path without updating*: The goal is to find the path with maximum reliability using mission prior knowledge only (Seshadri and Srinivasan, 2010).

4 *Most energy-efficient path with exploitation only*: In this approach, the path with the highest z-score is exploited. The distribution of energy requirement of road segments is updated using real-time measurements, but exploration is not utilised. The vehicle cannot change its course once undertaking a path unless the path is pruned by real-time measurements. If this situation does not occur, the UGV continues its course until it reaches another node.

5 *Most energy-efficient path with exploration and exploitation (proposed method)*: In this approach, we implement the methodology introduced in the previous section.

The prior distributions of road segments are obtained from experimental studies of Sadrpour et al. (2013a). For simulation, the power data is generated using the surrogate model introduced by Sadrpour et al. (2013b). The scaled aggressive EPA US06 driving cycle is used to represent the velocity profile of the UGV, over each segment as shown in Figure 5. The rolling resistance coefficients are generated using the normal distribution to capture the variations within the segment. The actual grade profiles were extracted from the Geocontext's (2013) study.

Figure 5 The EPA drive cycle was scaled both in time and speed magnitude in the simulation studies (see online version for colours)

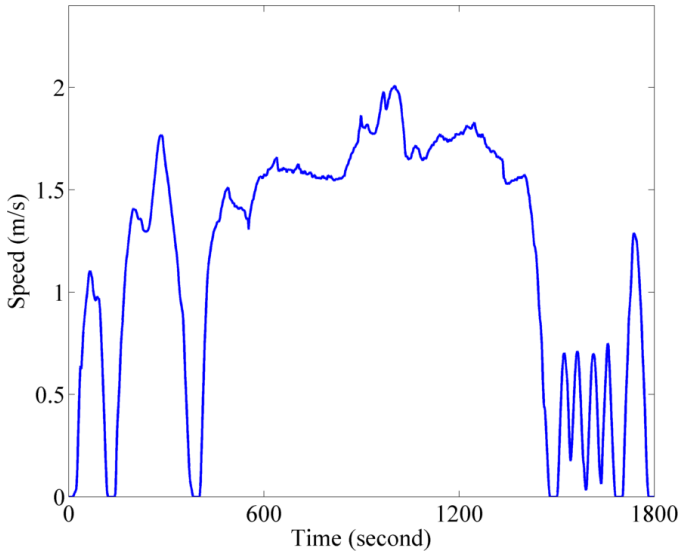

To validate the proposed approach, we demonstrate its application to a real-world scenario. In this scenario, the UGV traverses part of the University of Michigan's north campus, as shown in Figure 6a. The goal is to reach node 7 from node 1, using one of the five alternative paths listed in Figure 6a. The parameters of the study are listed in Table 1. Although the road segments' grade profiles were extracted from a database, it is assumed that this exact information is not generally available to the UGV operator. Figure 6b shows the schematic of the mission network with road segments' qualitative prior information, i.e. roads' average grades and surface conditions, and the length of each road, which are shown with numbers along each road segment. Road segment *d* is a new shortcut that the operator is not familiar with. The operator knows that the segment is sidewalk-(steep) uphill, and expresses their lack of knowledge by assigning a larger road grade prior variance (i.e. 1.5 times larger than a typical road grade variance). Also value of ξ shows if the operator is more interested in reducing the prediction uncertainty by exploration, or is more concerned about the battery remaining energy. Assigning a weight of $\xi = 0.1$ indicates that we are more interested in the reduction of uncertainty. A justification is that finding the minimum energy path is likely to save more energy in the long run, outweighing the potential energy savings from shorter exploration.

Figure 6 (a) The simulated case study network of alternative paths from node 1 to node 7. Five alternative paths have been identified as follows: $Q = \{q_1 = \{b, c, d, i\}, q_2 = \{a, e, f\}, q_3 = \{b, c, e, f\}, q_4 = \{a, d, i\}, q_5 = \{b, g, h, i\}\}$. (b) The mission network schematic with prior information about each road segment surface condition. The number along each road represents the length of the road in metres (see online version for colours)

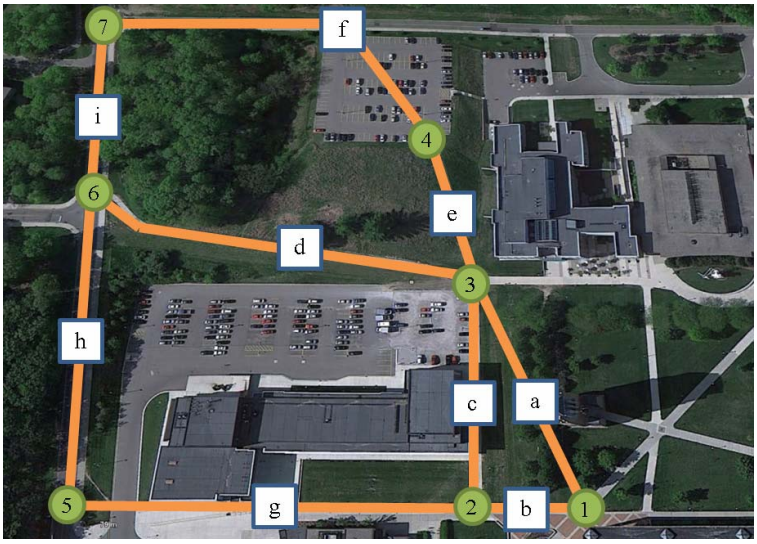

(a)

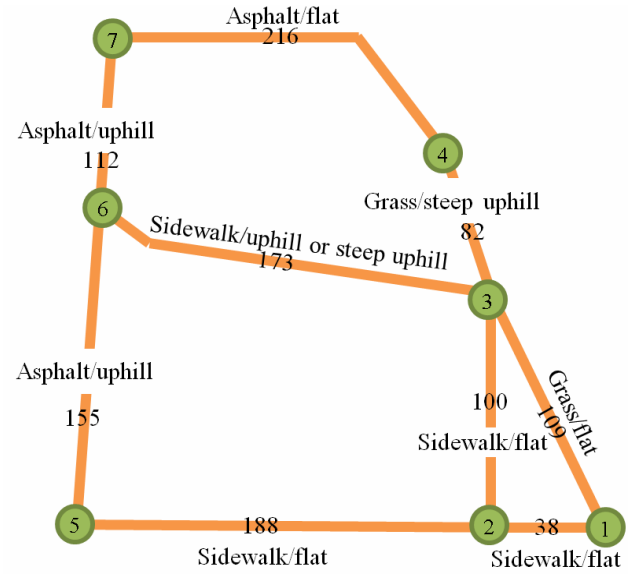

(b)

Table 1 Parameters of the simulated case study

m	40 kg	β	28 W	ξ	0.1
$C_{l'}$	$N(0.22, 0.003^2)$	σ_ε	7 W	Δt	1 s
z_α	2	T	85 KJ	γ	0.025
Examples of typical prior distributions of C					
Grass/flat	$N(0.319, 0.043^2)$		Asphalt/flat	$N(0.282, 0.044^2)$	
Sidewalk/flat	$N(0.276, 0.043^2)$		Grass/steep-uphill	$N(0.459, 0.058^2)$	
Asphalt/uphill	$N(0.352, 0.044^2)$		Sidewalk/uphill	$N(0.346, 0.043^2)$	

Figure 7a represents the initial distributions of energy requirements of paths based on the mission prior knowledge as well as the actual energy requirement of the paths along with the failure threshold. Using the naive approach, it is not clear which path should be selected since direct comparison of alternative paths is not conclusive. Using approaches 2 and 3, path q_2 has both the highest reliability and the lowest expected energy requirement. However, clearly, this path is not actually the most energy-efficient. Using approach 4 also, path q_2 is selected, and without exploration remains unpruned until the vehicle reaches the destination. Figure 7b depicts the updated energy requirement distribution of path q_2 using approach 4. While traversing road segments $\{e, f\}$ the predicted energy requirement generally has an increasing trend. The drop in the predicted energy at around observation 75 is because the initial part of road segment e is flat while the rest of it is steep-uphill and the operator states that the road is steep-uphill. The increase in the predicted energy at around observation 175 is due to slightly larger actual grade and rolling resistance compared to their prior distribution means.

Based on approach 5 if any of paths $\{q_1, q_2, q_3, q_4\}$ is selected, the UGV will explore all four road segments along them. Thus, the budget assignment optimisation needs to be solved at node 1 and 3 only. The number of measurements from $\{a, b\}$ are [4, 2]. The CPU time to perform the optimisation was approximately 2 s using a quad-core Intel core i-7 processor. After exploring road segments $\{a, b\}$, path q_2 still has the highest z-score. The UGV then exploits road segment a until it reaches node 3. The vehicle then explores the alternative road segments $\{e, d\}$ and collects [3, 4] measurements from each, respectively. The z-score associated with each path after the second exploration is shown in Figure 8a. Clearly, from Figure 7a, path $q_4 = \{a, d, i\}$, which actually has the least energy requirements (see Figure 6a for actual energy requirements), should be selected. Figure 8b depicts the predicted energy requirement of path q_4. The change in the decision from path q_2 to path q_4 after exploration is due to a drop in the predicted energy requirement of road segment d, shown with a circle in Figure 8b, which initially perceived to have an average grade of 6 degrees compared to an actual average grade of around 3 degrees uphill.

This case study shows the effectiveness and flexibility of the proposed approach. Operators can express their lack of knowledge of road surface conditions with large prior variances. These road segments will be assigned larger exploration budgets when solving the budget allocation optimisation problem. Also, we noticed that when generating alternative paths for different real-world scenarios, many alternative paths can be intuitively eliminated using the naive approach. Consequently, the number of alternative paths requiring exploration does not necessarily increase exponentially with the number of nodes and road segments.

Figure 7 (a) The initial energy requirement distribution of paths based on mission prior knowledge. The distributions are represented by their mean and confidence intervals. The actual energy requirement of each path, i.e. E_{qi}, is shown with a dashed line, and the energy failure threshold, i.e. T, is shown with a solid line. (b) The energy requirement of path q_2 vs. observations from this path using approach 4. The numbered circles correspond to the nodes in the network. This path cannot be pruned against the alternative paths and is traversed to reach to node 7 (see online version for colours)

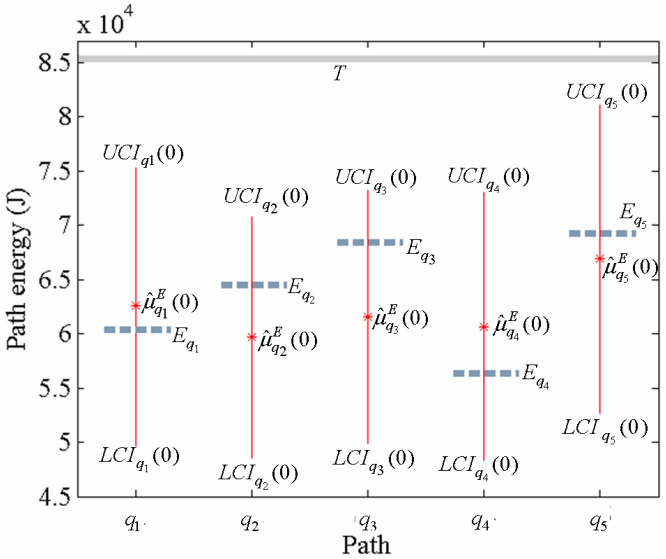

(a)

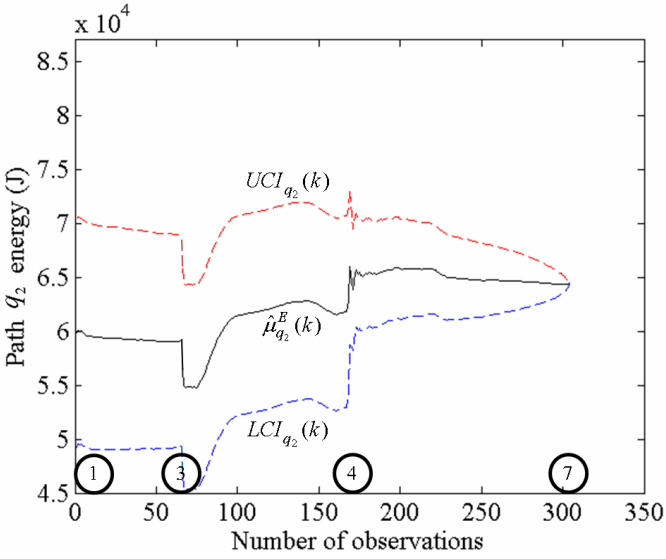

(b)

Figure 8 (a) Updated z-score of each path after exploring road segments $\{e, d\}$. The exploration stage requires a total of 13 measurements. Path $q_4 = \{a, d, i\}$ has the highest z-score and is subsequently selected to be traversed. Note that for example, path q_5 can be clearly eliminated at this stage due to very low z-score. (b) Path q_4 energy distribution updates using approach 5 (see online version for colours)

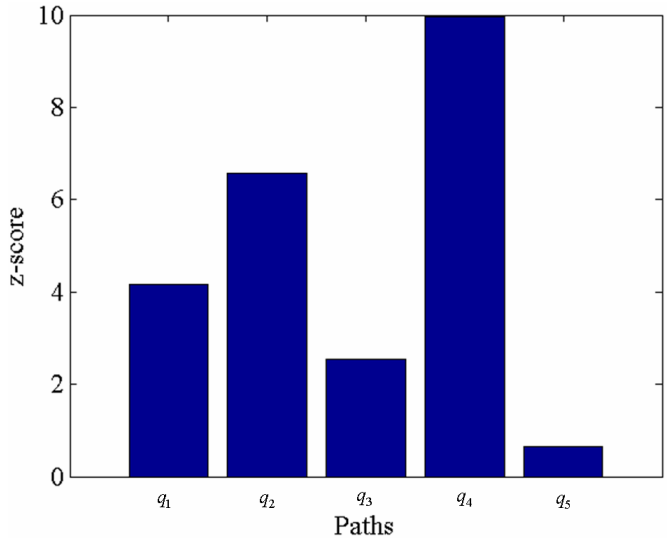

(a)

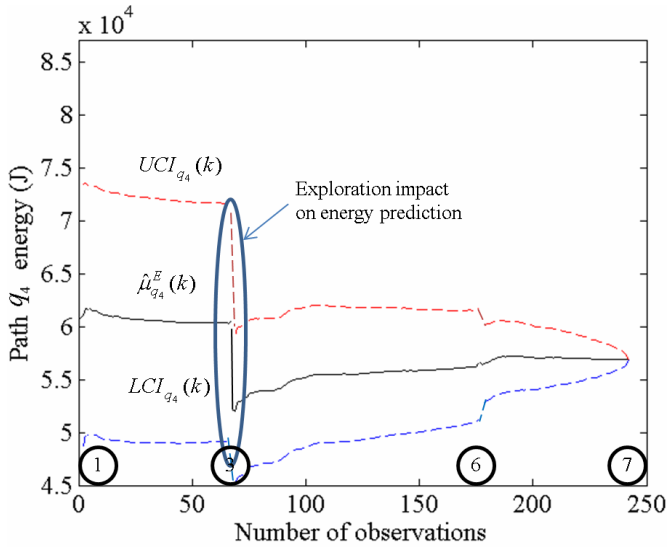

(b)

5 Concluding remarks

This paper proposes dynamic energy-efficient path planning for a UGV while minimising the probability of depleting the on-board energy during mission execution. Mission prior knowledge and real-time sensory information, i.e. instantaneous power consumption and velocity, are used to update the distribution of the predicted energy requirements of alternative paths through two decision-making stages. The first is an exploration stage, where prior information may not be sufficiently accurate, and an energy-efficient exploration strategy is used to reduce the uncertainty of the predicted energy requirements of the feasible paths. Also, paths that are unlikely to be the most reliable are removed through pruning. The second stage, i.e. exploitation, involves the UGV traversing the most reliable path based on the results of the first stage. Our simulated case study shows that the proposed approach outperforms several other potential offline and online path planning methods. Future research will focus on experimental validation studies and ways to reduce the computational complexity of the proposed approach for larger networks.

References

Applegate, D.L., Bixby, R.E., Cook, W.J. and Chvatal, V. (2011) *The Traveling Salesman Problem: A Computational Study*, Princeton University Press.

Borgstrom, P., Singh, A., Jordan, B., Sukhatme, G., Batalin, M. and Kaiser, W. (2008) 'Energy based path planning for a novel cabled robotic system', *International Conference on Intelligent Robots and Systems (IROS)*, pp.1745–1751.

Broderick, J., Tilbury, D. and Atkins, E. (2012) 'Energy usage for UGVs executing coverage tasks', *Proceedings of Unmanned Systems Technology*, Baltimore, Maryland, USA, pp.83871A–83871A.

Carlson, J. and Murphy, R. (2003) 'Reliability analysis of mobile robots', *International Conference on Robotics and Automation (ICRA)*, Vol. 1, Taipei, Taiwan, pp.274–281.

Carlson, J. and Murphy, R. (2005) 'How UGVs physically fail in the field', *IEEE Transactions on Robotics*, Vol. 21, No. 3, pp.423–437.

Carlson, J., Murphy, R. and Nelson, A. (2004) 'Follow-up analysis of mobile robot failures', *International Conference on Robotics and Automation (ICRA)*, New Orleans, Louisiana, USA, pp.4987–4994.

Congdon, P. (2003) *Bayesian statistical modeling*, John Wiley & Sons, Ltd.

Dearden, R., Friedman, N. and Andre, D. (1999) 'Model based Bayesian exploration', *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI)*, San Francisco, CA, USA, pp.150–159.

Dearden, R., Friedman, N. and Russell, S. (1998) 'Bayesian Q-learning', *Proceedings of Fifteenth National Conference on Artificial Intelligence (AAAI)*, Menlo Park, CA, USA, pp.761–768.

Denardo, E.V. (2003) *Dynamic Programming: Models and Applications*, Dover Publications, Englewood Cliffs, NJ.

Fan, Y.Y., Kalaba, R.E. and Moore, J.E. (2005) 'Shortest paths in stochastic networks with correlated link costs', *Computers and Mathematics with Applications*, Vol. 49, Nos. 9/10, pp.1549–1564.

Geocontext (2013) *Geocontext Center for Geographic Analysis*. Available online at: www.geocontext.org

Hoff, P.D. (2009) *A First Course in Bayesian Statistical Methods*, 1st ed., Springer.

Kramer, J. and Murphy, R. (2006) 'Endurance testing for safety, security and rescue robots', *Performance Metrics for Intelligent Systems*, Washington, DC, USA, pp.247–254.

Mei, Y., Lu, Y-H., Hu, Y.C. and Lee, C.S.G. (2004) 'Energy-efficient motion planning for mobile robots', *International Conference on Robotics and Automation (ICRA)*, Vol. 5, pp.4344–4349.

Mei, Y., Lu, Y-H., Hu, Y.C. and Lee, C.S.G. (2005) 'A case study of mobile robot's energy consumption and conservation techniques', *IEEE International Conference on Advanced Robotics*, Seattle, WA, USA, pp.492–497.

Powell, W.B. (2011) *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed., Wiley, Hoboken, NJ.

Ryzhov, I.O. and Powell, W.B. (2011) 'Information collection on a graph', *Operations Research*, Vol. 59, No. 1, pp.188–201.

Sadrpour, A., Jin, J. and Ulsoy, A.G. (2012) 'Mission energy prediction for unmanned ground vehicles', *International Conference on Robotics and Automation (ICRA)*, St. Paul, Minnesota, USA.

Sadrpour, A., Jin, J. and Ulsoy, A.G. (2013a) 'Experimental validation of mission energy prediction model for unmanned ground vehicles', *American Control Conference (ACC)*, Washington, DC.

Sadrpour, A., Jin, J. and Ulsoy, A.G. (2013b) 'Mission energy prediction for unmanned ground vehicles using real-time measurements and prior knowledge', *Journal of Field Robotics,* Vol. 30, No. 3, pp.399–414.

Sadrpour, A., Jin, J., Ulsoy, A.G. and Lee, H.J. (2011) 'Simulation-based acceptance testing for unmanned ground vehicles', *International Journal of Vehicle Autonomous Systems*, Vol. 11, No. 1, pp.62–85.

Seshadri, R. and Srinivasan, K.K. (2010) 'Algorithm for determining most reliable travel time path on network with normally distributed and correlated link travel times', *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2196, pp.83–92.

Sipahioglu, A., Yazici, A., Parlaktuna, O. and Gurel, U. (2008) 'Real-time tour construction for a mobile robot in a dynamic environment', *Robotics and Autonomous Systems*, Vol. 56, No. 4, pp.289–295.

Stancliff, S. and Dolan, J. (2005) *Towards a Predictive Model of Mobile Robot Reliability*, Technical Report.

Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, A Bradford Book, Cambridge, MA.

Tilbury, D.M. and Ulsoy, A.G. (2011) 'A new breed of robots that drive themselves', *ASME Mechanical Engineering Magazine*, Vol. 133, pp.28–33.

Ulsoy, A.G., Peng, H. and Cakmakci, M. (2012) *Automotive Control Systems*, Cambridge University Press, Cambridge, England.

Wei, H., Wang, B., Wang, Y., Shao, Z. and Chan, K.C.C. (2012) 'Staying-alive path planning with energy optimization for mobile robots', *Expert Systems with Applications*, Vol. 39, No. 3, pp.3559–3571.

Appendix

For ease of notation, we drop the index i that indicates the road segment whenever possible in Appendix.

A.1 Parameter estimation for random slope

Let us denote $\mathfrak{C}'_\ell = C + \mathfrak{C}_\ell$ where $\mathfrak{C}_\ell \sim N(0,\sigma^2)$ and thus $\mathfrak{C}'_\ell \sim N(C,\sigma^2)$. The estimation involves finding the posterior distribution of $\mathfrak{C}'_\ell$, C, and σ^2. The joint distribution of the data and parameters is proportional to

$$\begin{aligned}&\pi\left(\{\mathfrak{C}'_\ell\}_\ell, C, \sigma^2 \middle| \{y_{i\ell}\}_{i\ell}, \{x_{i\ell}\}_{i\ell}, \sigma_\varepsilon^2, \mu^C, \sigma^C, \rho, \kappa\right)\\ &\quad\propto\left(\prod_{\ell=1}^{L}\left(\prod_{i=1}^{n_\ell}\pi\left(y_{i\ell}|\mathfrak{C}'_\ell, x_{i\ell}\right)\right)\pi\left(\mathfrak{C}'_\ell|C,\sigma^2\right)\right)\pi\left(C|\mu^C,\sigma^C\right)\pi\left(\sigma^2|\rho,\kappa\right)\end{aligned} \tag{A1}$$

where $\pi\left(y_{i\ell}|\mathfrak{C}'_\ell, x_{i\ell}\right)$, $\pi\left(\mathfrak{C}'_\ell|C,\sigma^2\right)$, and $\pi\left(C|\mu^C,\sigma^C\right)$ are normal distributions, and $\pi\left(\sigma^2|\rho,\kappa\right)$ (i.e. the prior distribution of σ^2) is an inverse-gamma distribution, and $x_{i\ell}$, $y_{i\ell}$ are the measured input and output from sub-segment ℓ. Parameters ρ,κ can be chosen so that the mean and variance of prior distribution matches an estimated value obtained through offline experiments.

Since the posterior distribution does not have a closed-form solution, we will use Gibbs sampling to draw samples from the full conditional distributions of parameters. Let us assume that data from $\ell=1,\ldots,j'$ sub-segments within the road segment have been collected. The full conditional distribution of $\mathfrak{C}'_\ell$ and C are as follows (Hoff, 2009):

$$\begin{aligned}&\pi\left(\mathfrak{C}'_\ell|\bullet\right)\\ &\sim N\left(\left(\sigma^{-2}+\sigma_\varepsilon^{-2}\sum_{i=1}^{n_\ell}x_{i\ell}^2\right)^{-1}\left(\sigma^{-2}C+\sigma_\varepsilon^{-2}\sum_{i=1}^{n_\ell}x_{i\ell}y_{i\ell}\right),\left(\sigma^{-2}+\sigma_\varepsilon^{-2}\sum_{i=1}^{n_\ell}x_{i\ell}^2\right)^{-1}\right)\end{aligned} \tag{A2}$$

where $\bullet$ indicates that the distribution is conditioned on all the remaining parameters.

$$\begin{aligned}&\pi\left(C|\bullet\right)\\ &\sim N\left(\left(\sigma^{C^{-2}}+j'\sigma^{-2}\right)^{-1}\left(\sigma^{C^{-2}}\mu^C+\sigma^{-2}\sum_{\ell=1}^{j'}\mathfrak{C}'_\ell\right),\left(\sigma^{C^{-2}}+j'\sigma^{-2}\right)^{-1}\right)\end{aligned} \tag{A3}$$

And, the full conditional distribution of σ^2 is as follows:

$$\pi\left(\sigma^2|\bullet\right)\sim\text{Inverse-Gamma}\left(\rho+j'/2,,\kappa+\sum_{\ell=1}^{j'}\left(\mathfrak{C}'_\ell-C\right)^2/2\right) \tag{A4}$$

The CPU time to perform the Gibbs sampling varied between 3 and 12 s using a quad-core Intel core i-7 processor. The computation time depends on: (a) number of draws in the Gibbs sampling; (b) size of the sub-segments; and (c) size of the road.

A.2 Predictive distribution of output

Suppose the vehicle has collected k measurements from a road segment, and $\hat{x}(k+j|k)$, $\forall j \geq 1$ is the expected input at time $k+j$, which is estimated by the weighted average of velocity measurements up to time k. The mean and variance of the j-step-ahead prediction of output are estimated as follows:

$$\begin{aligned} E\left[y\left(k+j|k,\{\mathfrak{C}_\ell\}_\ell=0\right)\right] &= E\left[E\left[y\left(k+j|k,\hat{x}\left(k+j|k\right),C\left(k|k-1\right),\{\mathfrak{C}_\ell\}_\ell=0\right)\right]\right] \\ &= \hat{\mu}^C\left(k|k-1\right)\hat{x}\left(k+j|k\right) \end{aligned} \tag{A5}$$

We use the fixed effect model to estimate the expected energy requirement of a road segment in real-time. The variance of prediction for j step-ahead prediction of power, which is assumed to belong to sub-segment ℓ, at k is calculated as follows:

$$\begin{aligned} var\left(y\left(k+j|k\right)\right) &= E\left[var\left(y(k+j|k,\hat{x}(k+j),C(k|k-1),\mathfrak{C}_\ell\right)\right] \\ &\quad + var\left(E\left[y\left(k+j|k,\hat{x}\left(k+j|k\right),C\left(k|k-1\right),\mathfrak{C}_\ell\right)\right]\right) \\ &\approx \sigma_\varepsilon^2 + \left(\left[\hat{\sigma}^C(k|k-1)\right]^2 + E\left[\sigma^2|k\right]\right)\hat{x}^2\left(k+j|k\right) \end{aligned} \tag{A6}$$

where $C(k|k-1)$ is estimated in real-time using the fixed effect model and $E[\sigma^2|k]$ is posterior mean of σ^2, which is estimated less frequently using the mixed effect model.

The covariance of prediction error of j and j' step-ahead predictions if both belong to the same sub-segment ℓ is calculated as follows:

$$\begin{aligned} &cov\left(y(k+j|k),y(k+j'|k)\right) = E\left[cov\left(y(k+j|k,\hat{x}(k+j),C(k|k-1),\mathfrak{C}_\ell\right),\right. \\ &\left.y\left(k+j'|k,\hat{x}\left(k+j'|k\right),C\left(k|k-1\right),\mathfrak{C}_\ell\right)\right)\Big] \\ &+cov\left(E\left[y\left(k+j|k,\hat{x}\left(k+j|k\right),C\left(k|k-1\right),\mathfrak{C}_\ell\right)\right],\right. \\ &\left.E\left[y\left(k+j'|k,\hat{x}(k+i|k),C\left(k|k-1\right),\mathfrak{C}_\ell\right)\right]\right) \\ &\qquad \approx \left(\left[\hat{\sigma}^C\left(k|k-1\right)\right]^2 + E\left[\sigma^2|k\right]\right)\hat{x}\left(k+j|k\right)\hat{x}\left(k+j'|k\right) \end{aligned} \tag{A7}$$

if j and j' do not belong to the same sub-segment their covariance is estimated by

$$cov\left(y\left(k+j|k\right),y\left(k+j'|k\right)\right) = \left[\hat{\sigma}^C\left(k|k-1\right)\right]^2\hat{x}\left(k+j|k\right)\hat{x}\left(k+j'|k\right) \tag{A8}$$

Since at k, $\hat{x}(k+j)=\hat{x}(k+j')$ the last term in the equation above becomes $[\hat{\sigma}^C\left(k|k-1\right)\hat{x}\left(k+j|k\right)]^2$. In the work of Sadrpour et al. (2012), we showed power consumption can be estimated by $P\left(k+j|k\right)\approx y\left(k+j|k\right)+\beta$. The term $\hat{E}^r(k)$ in equation (4) is estimated as follows:

$$\hat{E}^r(k) \approx \sum_{j=1}^{n} P\left(k+j|k\right)\Delta t \tag{A9}$$

where n is the expected number of remaining measurements from the road segment and is estimated by $\hat{r}(k)/(\hat{u}(k)\Delta t)$ where $\hat{u}(k)$ is the weighted average of past velocity measurements. Let us assume the vehicle is at sub-segment ℓ', and let r_ℓ denote the remaining distance from sub-segment ℓ. Using equation (A5)–(A8), the expected value and variance of $\hat{E}^r(k)$ are as follows:

$$\begin{aligned} E\left[\hat{E}^r(k)\right] &= \left(\hat{\mu}^C(k|k-1)\hat{x}(k+j)+\beta\right)\left(\hat{r}(k)/(\hat{u}(k)\Delta t)\right)\Delta t \\ &= \hat{r}(k)W\hat{\mu}^C\left(k|k-1\right)+\hat{t}(k)\beta \end{aligned} \tag{A10}$$

$$\begin{aligned} var\left(\hat{E}^r(k)\right) &= var\left(\sum_{j=1}^{n} P\left(k+j|k\right)\right)\Delta t^2 = W^2\hat{r}^2(k)\left[\sigma^C\left(k|k-1\right)\right]^2 \\ &+\hat{t}\left(k\right)\hat{\sigma}_\varepsilon^2\Delta t+W^2E\left[\sigma^2|k\right]\sum_{\ell=\ell'}^{L} r_\ell^{\,2} \end{aligned} \tag{A11}$$

If we assume that each sub-segment has roughly an equal length r_c, we can simplify the term $W^2E\left[\sigma^2|k\right]\ \sum_{\ell=\ell'}^{L} r_\ell^{\,2}$ to $W^2E\left[\sigma^2|k\right]r_c\hat{r}\left(k\right)$ where $W^2E\left[\sigma^2|k\right]r_c$ is a constant denoted by ϕ.

A.3 Posterior updating using similarities

Consider the following distribution for combined parameters of two road segments, i,i', $\pi\left(C_i,C_{i'}\right)\sim N(\mu,\Sigma)$. Using realisations from C_i, the distribution of $C_{i'}$ can be updated by the conditional distribution $\pi\left(C_{i'}|C_i\right)$. However, realisations of combined parameter $C_{i'}$ are not available. The only measurable quantities are the input and output of model [(equation (2)]. Let us assume k measurements are collected from road segment i denoted by obs_i. The distribution of C_i is updated as follows:

$$\pi\left(C_{i'}|obs_i\right)=\int\pi\left(C_{i'}|C_i\right)\pi\left(C_i|obs_i\right)dC_i \tag{A12}$$

where $\pi\left(C_i|obs_i\right)$ is the posterior distribution of C_i. While this updating scheme can be applied to any multivariate normal distribution, in our case, we assume that two similar roads share the same combined parameter having a correlation of one. Thus, any realisation of C_i from $\pi\left(C_i|obs_i\right)$ is a realisation from C_i, and based on equation (A12), the combined parameter of both road segments can be simultaneously updated using measurements from one of them.

A.4 Covariance of energy requirement of road segments

Two road segments are similar if they share the same C. Let us assume k measurements have been collected from road segment i and the vehicle has not yet started to traverse similar road segment i'. Two predictions from road segments i and i' has a covariance

of $\left[\hat{\sigma}^C(k|k-1)\right]^2 \hat{x}\left(k+j|k\right)^2$. Thus, covariance of energy prediction can be calculated using the derivations from Section A.2 as follows:

$$\begin{aligned} &cov\left(\sum_{j=1}^{n_i-k} y\left(k+j|k\right)\Delta t, \sum_{j'=1}^{n_{i'}} y\left(j'|k\right)\Delta t\right) \\ &=\left(n_i-k\right)n_{i'}\left[\hat{\sigma}^C\left(k|k-1\right)\right]^2 \hat{x}\left(k+j|k\right)^2 \Delta t^2 = W^2 \hat{r}_i(k)\hat{r}_{i'}(k)\left[\hat{\sigma}^C\left(k|k-1\right)\right]^2 \end{aligned} \tag{A13}$$

where $n_i, n_{i'}$ are the number of measurements from road segments *i* and i', respectively.

A.5 Estimating the reduction in uncertainty

For estimating the reduction of energy prediction variance, we use the prediction variance relation in equation (7). We assume based on prior knowledge of driving style, the velocity follows a normal distribution $u \sim NID\left(\mu_u, \sigma_u^2\right)$. The posterior distribution of *C* after *k* measurements is estimated by $\left[\hat{\sigma}^C\left(k\right)\right]^2 = \left(\left[\hat{\sigma}^C\left(0\right)\right]^{-2} + \sigma_\varepsilon^{-2}\sum_{i=1}^{k} x^2\left(i\right)\right)^{-1}$. The expected value of $E\left[\sum_{i=1}^{k} x^2(i)\right] = kW^2\left(\mu_u{}^2 + \sigma_u^2\right)$. The two other parameters of equation (7) $\hat{r}(k)$ and $\hat{t}(k)$ are estimated by $\left(d - \mu_u k\Delta t\right)$ and $\left(d/\mu_u - k\Delta t\right)$, respectively. Replacing these estimates in equation (7), we obtain equation (11).