

Enhanced Stochastic Mobility Prediction on Unstructured Terrain using Multi-Output Gaussian Processes

Sin Ting Lui¹, Thierry Peynot^{1,2}, Robert Fitch¹, and Salah Sukkarieh¹

¹ Australian Centre for Field Robotics, The University of Sydney, Sydney, Australia,
{a.lui,rfitc,alah}@acfr.usyd.edu.au

² School of Electrical Engineering and Computer Science, Queensland University of
Technology, Brisbane, Australia, t.peynot@qut.edu.au

Abstract. Outdoor robots such as planetary rovers must be able to navigate safely and reliably in order to successfully perform missions in remote or hostile environments. Mobility prediction is critical to achieving this goal due to the inherent control uncertainty faced by robots traversing natural terrain. We propose a novel algorithm for stochastic mobility prediction based on multi-output Gaussian process regression. Our algorithm considers the correlation between heading and distance uncertainty and provides a predictive model that can easily be exploited by motion planning algorithms. We evaluate our method experimentally and report results from over 30 trials in a Mars-analogue environment that demonstrate the effectiveness of our method and illustrate the importance of mobility prediction in navigating challenging terrain.

Keywords: mobility prediction, learning, planetary rover, stochastic motion planning

1 Introduction

Safe and reliable navigation for outdoor robots involves principled consideration of control uncertainty. We are interested in *mobility prediction* for systems such as planetary rovers that are mechanically designed for mobility in natural environments, but which pose interesting challenges in planning and control due to complex interactions with rigid or deformable terrain. Addressing these challenges is important in enabling the application of planetary rovers that operate outside the possibility of human intervention, especially since the most interesting scientific tasks must often be performed in the most difficult terrain [1]. In recent work [2, 3], we explored mobility prediction using Gaussian process (GP) regression models where heading (bearing) and distance uncertainty are treated independently. Here we consider the case where heading distance uncertainty are coupled, and model this coupling using a multi-output GP.

Mobility prediction is the problem of estimating the likely behaviour of the rover in response to a given control action on given terrain. The goal is to provide a predictive model of control uncertainty that can be exploited by planning algorithms to find *safe paths*. This idea is different from classical motion planning, which seeks to minimise

time or distance while avoiding obstacles. Safe paths, in contrast, have low likelihood of leading the robot into unsafe states during execution, such as tipping over and getting stuck, in addition to collisions. On unstructured terrain, mobility prediction is difficult due to the complex terramechanics involved [4], and is distinct from the large body of work in terrain classification [5, 6] which seeks to detect and avoid hazards but does not attempt to build a predictive model of control uncertainty.

In our recent work [3] we proposed a mobility prediction method that learns a stochastic transition model from previous experience. This method considers the effects of terrain interaction on the macroscopic behaviour of the rover without modelling detailed wheel-soil interactions. We demonstrated the effectiveness of this method experimentally in Mars-analogue terrain where path executions were safer and more reliable in rigid and deformable terrain than paths generated by a classical motion planning algorithm. The model consists of multiple single-output GP components that predict heading and distance uncertainty independently.

In this paper, we propose a new mobility prediction method that captures the correlation between heading and distance uncertainty. We model correlated heading and distance using multi-output GP regression, where the outputs represent the expected resulting position of the rover with respect to a given control action. We also propose a new representation of terrain features that improves on our previous method by considering a larger area of terrain in the intended direction of motion.

We present experimental results that show significant improvement over the single-output GP method, and compare both results to the no-uncertainty case as a control condition. First, we evaluate the mobility prediction model in isolation and discuss terrain types where correlated heading and distance uncertainty is beneficial. Then, we present results from over 30 navigation trials. All experiments were performed using a six-wheeled planetary rover platform in challenging Mars-analogue terrain. In the control condition (no uncertainty considered), no trials could be completed successfully due to rocks and deformable terrain. All trials were successful in the multi-output GP condition, only one trial failed in the single-output condition, and the multi-output case resulted in demonstrably safer paths.

The paper is organised as follows. Section 2 discusses related work. Section 3 presents our general mobility prediction approach as background material for completeness. In Secs. 4 and 5 we present our new algorithms and implementation. Experimental results are reported in Secs. 6 and 7, and Sec. 8 concludes the paper.

2 Related work

Robust terrain traversability estimation and navigation is an important topic of research, especially in the context of planetary exploration. *Traversability analysis* is the general problem of assessing to what degree a robot may traverse given terrain [7]. Typically this analysis is performed by examining local terrain geometry [8] and soil properties [9]. *Terramechanics* is the study of wheel-soil interactions [10], and is difficult to apply online due to large parameter uncertainty even in homogeneous terrain [11]. *Near-to-far learning* is an online terrain classification approach where the association between proprioceptive measurements and the corresponding classes allow remote ter-

rain to be classified based on the rover’s previous experience [12, 5]. Related work in stochastic mobility prediction typically seeks to reactively compensate for control uncertainty due to factors such as slip, but assumes a reference path [13, 12].

Our approach is data-driven and relies on Gaussian process regression, a machine learning technique that has recently gained popularity in robotics applications. GPs are non-parametric, do not assume an underlying function shape, and provide a continuous estimate of prediction uncertainty [14]. GPs have natural application to spatially correlated data with sparse datasets. Multi-output GPs, also known as dependent GPs, allow correlated outputs to be simultaneously learnt [15]. In the context of regression, several implementations of multi-output GPs have been proposed [16, 17]. To the best of our knowledge, multi-output GPs have not previously been applied and experimentally validated in the context of mobility prediction.

3 Background

In prior work [3], we proposed a path planning approach that accounts for control uncertainty by learning a Stochastic Mobility Prediction Model (SMPM) from experience. In this paper, we build on this previous work by tackling two main limitations: 1) considering the control uncertainty in multiple dimensions jointly, rather than each dimension separately, and 2) taking into account the variability of the paths taken by the rover when collecting the input information about the terrain. In this section, we summarise the original technique for planning with control uncertainty using a learned SMPM.

3.1 Stochastic Mobility Prediction Model

Given a rover state $s = \{x, y, \psi\}$, where (x, y) is the 2D position of the rover and ψ its orientation (yaw), the execution of a given action $a \in A$ will result in state s' . A deterministic mobility prediction model is commonly used to represent the transition from s to s' . However, due to control uncertainty, in practice the resultant state is not deterministic. This can be accounted for by formulating the state transition function as a probability density function of the relative transition between states, $p(\Delta s|s, a)$, with $\Delta s \equiv s' - s$. In this work, Δs is defined using a polar representation of the state space. The N components Δs_i of Δs are the vehicle’s heading, distance and yaw:

$$\begin{aligned} \Delta s &\triangleq \{\Delta s_{head}, \Delta s_{dist}, \Delta s_{yaw}\} \\ &\triangleq \{\tan^{-1}(\Delta y, \Delta x), \sqrt{(\Delta x)^2 + (\Delta y)^2}, \Delta \psi\}. \end{aligned}$$

Since the outcome of an action is strongly correlated with the geometry of the unstructured terrain the rover has to traverse, the transition model depends on terrain profiles $\lambda(s, a)$, which contain information on terrain geometry. In our prior work, $\lambda(s, a)$ encoded the variations of rover’s attitude and configuration angles experienced between s and s' , predicted using a kinematic model (see [3]). Since the exact path between s and s' is not known in advance, these predictions were made at discrete locations along a straight line drawn between the initial state s and the average resultant state $\bar{s}' = s + \Delta s_a$ for this action in the training data (see Sec. 4.2 for more details).

To learn the SMPM from experience, training data were collected in a representative environment by performing multiple executions of each action a over a variety of terrain profiles. Since training can only provide a limited, sampled subset of the feature space, the SMPM was then learnt using Gaussian Process regression. This consists in learning the correlations K between the outcomes of each action a and the corresponding terrain profiles $\lambda(s, a)$. Once this training is complete, given an action a , we can query the SMPM for a prediction of the expected control error distribution for any terrain profile $\lambda_*(s, a)$ on similar terrain. For each action $a \in A$, the distribution can be written as:

$$p(\Delta s_a - \overline{\Delta s}_a | \lambda(s, a), a), \quad (1)$$

where $\overline{\Delta s}_a$ is the mean value of Δs_a across all executions of action a in the training data. $\Delta s_a - \overline{\Delta s}_a$ represents the discrepancy between the actual execution and the expected action execution (based on the raw training data). Note that with this formulation the training data for each action has zero mean.

3.2 Single-Output GP Learning

For each action and component Δs_i of Δs , given a training set of n input features $X = \{x_j | j = 1, \dots, n\}$ and their corresponding action outcomes, or *targets* $Z = \{z_j\}$, the GP can provide a predictive distribution g_* for any query inputs x_* [14]. g_* is estimated as the Gaussian distribution:

$$p(g_* | X, Z, x_*) \sim \mathcal{N}(\mu_*, \Sigma_*), \quad (2)$$

with predictive mean

$$\mu_* = K(x_*, X)[K(X, X) + \sigma_n^2 I]^{-1} Z, \quad (3)$$

and variance

$$\Sigma_* = K(x_*, x_*) - K(x_*, X)[K(X, X) + \sigma_n^2 I]^{-1} K(X, x_*), \quad (4)$$

where $K(X, x_*)$ is the covariance function that describes the spatial correlation between two inputs X and x_* , I is the identity matrix, and σ_n is the noise variance.

In our prior work each component of Δs_i , $i \in \llbracket 1, N \rrbracket$ for each action $a \in A$ is estimated by a different GP. The predictive distribution of each GP can be written as:

$$p(g_* | X, Z, x_*) = p(\Delta s_{i,a} - \overline{\Delta s}_{i,a} | \lambda(s, a), a) \sim \mathcal{N}(\mu_*, \Sigma_*), \quad (5)$$

where $\Delta s_{i,a}$ is the i^{th} component of the change of state Δs resulting from executing action $a \in A$, and $\overline{\Delta s}_{i,a}$ is the mean value of Δs_i across all executions of action a in the training data. In this implementation of mobility prediction, a training input (or a query input x_*) is a terrain profile $x_j = \lambda(s, a)$, and a target is the corresponding action outcome: $z = \Delta s_{i,a} - \overline{\Delta s}_{i,a}$. The uncertainty in each component Δs_i is accounted for by using the full distribution learned from Δs_i and the expectation of the other components. We then use the learnt SMPM as a transition model during planning (see Fig. 1).

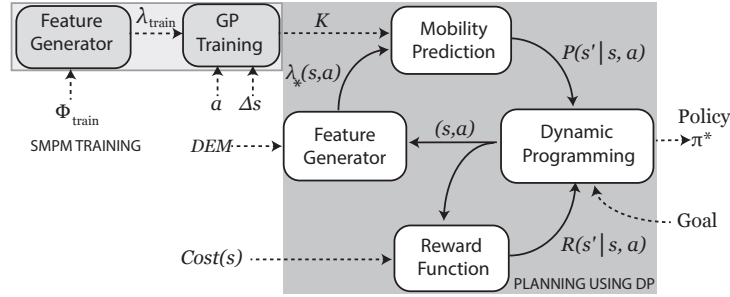


Fig. 1: System outline. The left box shows the training conducted to build the mobility prediction model, while the planning is shown in the box to the right. During the training stage, training terrain profiles λ_{train} are generated from sets of $\{\Phi_{train}\}$, representing the attitude and configuration Φ_{train} of the platform evaluated at regular intervals along the traversed terrain trajectory when executing action a (see Fig. 2c). This training process produces K , which is used to estimate continuous SMPPM with a GP. Once training is complete, $\lambda_*(s, a)$ terrain profiles are generated from the DEM of the terrain that the rover needs to traverse. These are then used to compute the stochastic transition distribution $P(s' | s, a)$. Given $P(s' | s, a)$ and the reward function $R(s' | s, a)$ for the terrain to be traversed, DP generates an optimal policy π^* .

3.3 Planning

Various planning methods can be used to exploit our mobility prediction model. We use a Markov decision process (MDP) formulation of the problem, where the transition function $P(s' | s, a)$ is provided by the SMPPM, and the reward is a function of action cost and vehicle's safety over the terrain (represented by a cost $Cost(s)$). Prior to planning, a cost map is computed using kinematic modelling predictions on a digital elevation map (DEM) generated using exteroceptive sensors on the rover. We then compute policies using dynamic programming to maximise the sum of rewards accumulated over sequences of actions [18]. In operation, the rover follows the policy: at the end of each action execution, the policy provides the rover with the next most appropriate action that it should execute from its current location.

The experiments conducted in [3] considered uncertainty in heading and in distance travelled, independently. The results indicated that by using our learned SMPPM to consider control uncertainty in the planning stage, at the execution we obtained paths with significantly reduced cost, i.e. safer and more efficient paths. Besides, the impact was stronger when considering heading uncertainty rather than distance uncertainty. Figure 1 gives an outline of the implementation of this approach.

4 Enhanced Mobility Prediction Modelling

In this paper, we propose to address two of the main limitations of the previous work. First, we propose to learn the multiple dimensions of control uncertainty jointly instead of independently. Second, we enhance the strategy used to collect the features representing the appropriate terrain profiles.

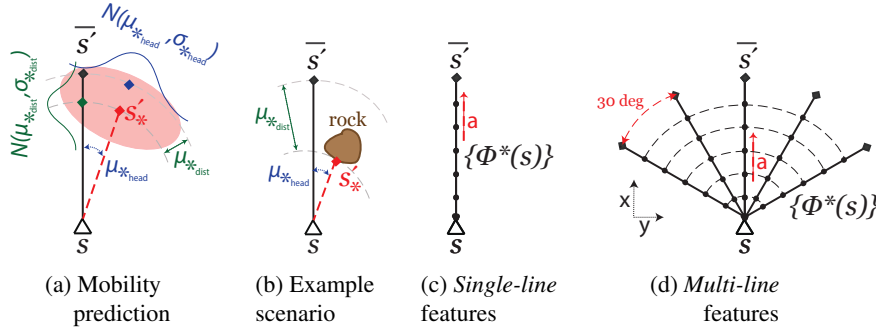


Fig. 2: (a) Example of predictions of the end position (s') of the rover (triangle) after executing action a , starting from s . The predictions made by the prior approach are shown: in blue when considering uncertainty in heading only (diamond: best guess, Gaussian: predicted distribution of heading uncertainty) and in green when considering uncertainty in distance only. The prediction made by the new approach, which accounts for the correlation between heading and distance uncertainty, are shown in red, with the ellipse representing 2 standard deviations in both directions. (b) Illustration of the correlation between heading and distance uncertainty. (c) *Single-line* features strategy used in prior work. The platform configuration, Φ , is evaluated at regular intervals on a line between s and s' to form the set of $\{\Phi_{*}(s)\}$. (d) *Multi-line* features used in the new approach proposed in this paper. The Φ samples are collected along multiple lines to better reflect the variety of possible outcomes of the action.

4.1 Joint Predictions and Multi-Output Learning

In the prior approach, the SMPM represented the uncertainty in each dimension of Δs separately, as illustrated by the blue and green diamonds in Fig 2a. As a result, when considering the outcome of a given action for a given terrain profile, one dimension was considered uncertain while the other was considered deterministic. For example, when accounting for uncertainty in heading, the distance that the rover was expected to travel during the execution of action a was assumed to be $\bar{\Delta s}_{dist,a}$, i.e. the average distance travelled by the rover during all executions of action a in the training data. However, in practice, the outputs of the prediction process (heading and distance deviations) may be highly correlated.

Consider the following example. The rover plans to execute the action of going straight ahead (see Fig. 2b). The terrain is sandy but flat, except for a rock located on the right hand side of its course, far enough that the rover would not touch it when driving perfectly straight. The rock is most likely not traversable by the vehicle. In practice, during the action execution, sometimes the rover will deviate on the right and get stuck against the rock. In such case, it is clear that the heading deviation experienced by the rover has a strong impact on the distribution of distance traveled by the rover, illustrating the correlation between heading and distance uncertainty in mobility prediction.

Figure 2a shows an example of prediction obtained when considering uncertainty in both heading and distance, and their correlation (in red), compared with the predictions made when considering only one dimension of uncertainty (in blue and green). To

address this issue, in this paper we propose to use multi-output GPs to learn the joint effects of control uncertainty.

Joint predictions of the correlated outputs are possible by using multi-output learning, however, defining the covariance matrix K can be difficult while guaranteeing its positive-definitiveness, required for GP regression. To model the correlated outputs, one method is to utilise Convolution Processes [17] between a smoothing kernel k_q and a latent function $u(z)$. The set of Q functions, representing the N correlated outputs, can be written as:

$$f_q(x) = k_q(x) * u(x) = \int_{-\infty}^{\infty} k_q(x-z)u(z)dz \quad (6)$$

where x is the input and z is the output (target).

In our approach, the smoothing kernel k_q used is the squared exponential function with heteroscedastic noise:

$$k_q(x-z) = \frac{S_q|M_q|^{1/2}}{(2\pi)^{p/2}} \exp\left[-\frac{1}{2}(x-z)^T M_q(x-z)\right], \quad (7)$$

where M_q, S_q and p are hyperparameters of the kernel. This kernel produces very smooth functions [14], can be integrated against most functions, and is widely used within the GP learning community.

The influence of noise $w_q(x)$ and R latent functions are considered on the function y_q by assuming that each output is independently corrupted. The function $y_q(x)$ can be expressed as:

$$y_q(x) = f_q(x) + w_q(x) = \sum_{r=1}^R \int_{-\infty}^{\infty} k_{qr}(x-z)u_r(z)dz + w_q(x). \quad (8)$$

The Σ_* and μ_* values required for generating the predictive distribution (see Eq. (5)) can be computed from $y_q(x_*)$. To consider the influence of multiple latent functions on $y_q(x)$, each latent function is assumed to be an independent GP. The covariance between two functions $y_q(x)$ and $y_s(x)$ can therefore be written as:

$$\text{cov}[f_q(x), f_s(x')] = \sum_{r=1}^R \int_{-\infty}^{\infty} k_{qr}(x-z) \int_{-\infty}^{\infty} k_{sr}(x'-z')k_{u_r u_r}(z, z')dz' dz, \quad (9)$$

where k_{qr} and k_{sr} are the kernel functions for the latent functions, and $k_{u_r u_r}$ is the covariance function for $u_r(z)$. The correlation between any given output $f_q(x)$ and the latent function $u_r(z)$ can be computed as:

$$\text{cov}[f_q(x), u_r(z)] = \int_{-\infty}^{\infty} k_{qr}(x-z')k_{u_r u_r}(z', z)dz'. \quad (10)$$

Joint predictions that consider more than one output can be calculated by using Eqs. (9) and (10), as the cross-covariance terms are incorporated into the estimation process. This enables us to consider the correlated heading and distance uncertainties jointly in the prediction of Σ_* and μ_* . In practice, for every action $a \in A$ the inputs

are the same as in the single-output GP formulation, however, a target represents the N -dimensional action outcome:

$$z = [\Delta s_{head} - \overline{\Delta s_{head}}, \Delta s_{dist} - \overline{\Delta s_{dist}}, \Delta s_{yaw} - \overline{\Delta s_{yaw}}]. \quad (11)$$

Furthermore, this means that for every action $a \in A$ we generate only one multi-output GP, as the predictive distribution of the multi-output GP estimates the N dimensions of Δs_i simultaneously. In this paper, the predictive distributions estimated by the multi-output GP over the entire state space constitute the SMPM employed in planning.

4.2 Terrain Profiles and Feature Generation

As mentioned previously, the estimation of the outcome of a given action a requires information about the profile of the terrain that the rover is going to traverse during the execution of this action (i.e. between the initial state s and the final state s'). In prior work, this information, which described the variations of the rover’s attitude and configuration angles (Φ), was collected at regularly-spaced discrete positions on a single straight line in the direction of the initial best guess of the outcome ($\overline{s'}$), see Fig. 2c. However, the prediction of the distribution of end states s'_* (see the ellipse in Fig. 2a) shows that the rover may actually follow a path that is quite different from this single line s to $\overline{s'}$, thereby travelling over different terrain geometry than anticipated initially.

Therefore, in this paper, we propose to expand the locations where we collect the relevant terrain geometry information (see Fig. 2d). We collect this information over multiple lines (five in practice), all starting from the initial state s , and placed at regular angular increments around the original single line, whose extremity is $\overline{s'}$ (we use increments of 30° in this paper). In order to use a fixed strategy for all queries of action outcomes, the extent of that coverage of the terrain was chosen to represent three standard deviations of the heading uncertainty observed in the training data.

In the next section we will show that features generated using these multiple lines better capture the possible configurations that the rover may encounter during the execution of each action. For convenience, in the remainder of the paper, we will refer to this new strategy to generate the features representing the appropriate terrain profiles as *multi-line features*, in contrast with the former approach, which uses *single-line features*, i.e. features captured over a single straight line.

5 Implementation

This section describes the implementation details of our proposed approach. The platform used in our experimental validation is *Mawson*, a six-wheeled holonomic rover prototype with a Rocker-bogie chassis, shown in Fig. 3. It is equipped with two visual cameras and an RGB-D camera (Microsoft Kinect), used as depth sensor only, mounted on a mast tilted down 20° for DEM generation. Other onboard sensors include three potentiometers to measure the bogie angles and the rocker differential, and an Intersense IS-1200 VisTracker device, comprising of an inertial measurement unit (IMU) and a camera, to compute 6-DOF rover localisation with $2cm$ average accuracy.

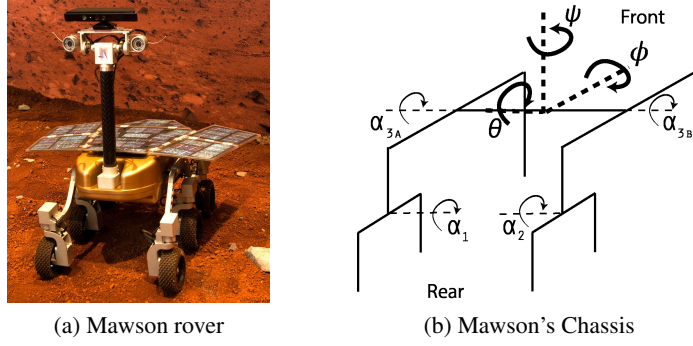


Fig. 3: The Mawson rover (a) and its attitude and configuration angles $\Phi = \{\phi, \theta, \alpha_1, \alpha_2\}$, shown in (b). ψ represents the yaw of the rover.

We define the same set of primitive actions that the rover can execute as in [3]: 8 crabbing actions, and 2 rotational actions, where crabbing is set at intervals of $\pi/4$ for a distance of about $0.3m$, and rotation is set at $\pm\pi/4$. These actions are initially calibrated on flat terrain, therefore, deviations from the original objective (both in heading and distance) are expected to happen in practice on rough terrain.

The experimental environment is a Mars-analogue terrain facility located inside the Sydney Powerhouse Museum in Australia. The terrain considered consists of both solid and loose soil, slopes of varying degrees, rocks of different sizes and shapes.

The DEMs used in this study are generated with a resolution of $0.05m \times 0.05m$ from point clouds acquired by the depth sensor. Kinematic predictions of attitude and internal configuration angles Φ were computed using a method similar to [19], allowing for the computation of features for the terrain profiles as well as cost maps.

As in [3], we define a terrain traversability cost as a function of the attitude and rocker bogie angles:

$$cost_{terrain}(s) = (\phi^2 + \theta^2 + 0.5(\alpha_1^2 - \alpha_2^2))^2, \quad (12)$$

where ϕ and θ are the roll and pitch of the platform, respectively, and (α_1, α_2) are internal angles of the rover's chassis (see Fig. 3b). This cost function captures the magnitudes of the platform's attitude and configuration during the traversal of the terrain, which are indicative of the difficulty and risk for the platform to traverse this patch of terrain. The reward function $R(s|s', a)$ used to compute the policies is defined as the average cost of states between the start state s and the resultant state s' , plus an action execution penalty ξ [3]:

$$R(s'|s, a) = -\xi - \frac{1}{M} \sum_{i=0}^M cost_{terrain} \left(s_x + \frac{i}{M}(s'_x - s_x), s_y + \frac{i}{M}(s'_y - s_y), s_\psi + \frac{i}{M}(s'_\psi - s_\psi) \right), \quad (13)$$

where $M = 20$ is the sampling resolution of the path between s and s' , and $\xi = 0.003$ is the penalty used in our implementation. During policy execution, the potentiometers

Table 1: Training data statistics. Mean values of the Δs_i components, by action.

Action:	<i>crab</i>					<i>rotate</i>
	(0π)	$(\pm\pi/4)$	$(\pm\pi/2)$	$(\pm3\pi/4)$	(π)	$(\pm\pi/4)$
# samples	72	118	139	114	72	100
Δs_{head} (rad)	-0.066	0.679	1.316	2.411	2.494	-
Δs_{dist} (m)	0.231	0.246	0.182	0.256	0.249	-
Δs_{yaw} (rad)	-	-	-	-	-	0.538

and the IMU measurements allow the attitude and internal configuration angles of the rover to be collected, such that terrain costs integrated over actual executed paths can be computed for the experimental analysis (see Sec. 7).

The set of most informative features to describe the terrain profiles, which are used for training and querying the GPs, were selected by performing a Principal Component Analysis (PCA) over a large variety of features capturing absolute values and variations within the Φ sets (see [3] for more details). These features, set in the vector λ , were collected on the terrain over multiple lines (for the proposed approach), or over a single straight line (for comparison with prior work), as described in Sec. 4.

To collect the required training data, the rover executed each action $a \in A$ multiple times over varying terrain profiles, while recording: the action a , the difference Δs between the end state of the rover after action execution s' and the start state s , and the platform attitude and configuration angles (Φ_{train}) during the action execution. Then, in order to train the GP, the features λ_{train} were systematically computed from Φ_{train} .

6 Experimental Validation of the Learned Mobility Prediction Model

In this section we validate the learned mobility prediction model experimentally, specifically demonstrating the benefits of 1) the multi-output GP learning, and 2) the extended feature set to better describe the relevant terrain profiles.

6.1 Training Data

We used the Mawson rover (see Sec. 5) and the training approach described in Sec. 3 to collect training data from more than 600 action executions, over numerous terrain profiles, varying from flat surface to rough terrain with significant slope.

Table 1 shows a summary of the training data obtained over *all* terrain for each action a . Note that due to the left-right symmetry of the platform, the training data were combined for symmetric actions. Therefore, only 6 different actions are shown in the table.

The table shows the average of the Δs_i components obtained over all executions of each action, on *all* terrain profiles experienced in the training phase. It can be noted that since the ability of the rocker-bogie chassis to overcome rough terrain depends on the rover orientation, the mean of Δs_i can be quite different for each action.

6.2 Mobility Prediction Model Validation

To cross-validate the proposed approach, we learned the mobility prediction models using 2/3 of the collected data, and tested the models using the remaining 1/3 of the data. Firstly, we show the benefits of multi-output learning compared with the state-of-the-art technique that uses single-output GPs. Secondly, we validate the use of extended features to capture the relevant terrain information as input of the mobility prediction.

Multi-Output Learning Results Table 2 shows the results obtained when predicting the outcomes of action a using multi-output GPs with the same strategy to collect information on terrain profiles as in prior work (i.e. *single-line features*), compared with the predictions made using the state-of-the-art approach with single-output GPs (which considers uncertainty in heading or distance, respectively). The *position errors* in the table correspond to the distance between the predicted end position of the rover and the actual end position given by the ground truth (i.e. the localisation system onboard the rover). The table provides the mean and standard deviation (std) of these position errors computed over all executions of each action in the test data. It can be seen that in all cases the results of the approach proposed in this paper are more accurate and more consistent than those obtained with the single-output GP that considers uncertainty in heading only. Compared with the single-output GP that considers uncertainty in distance only, the results are comparably accurate for the first three actions, and more accurate for $crab(\pm 3\pi/4)$ and $crab(\pm\pi)$. The benefit is particularly significant for the latter. Figure 4 illustrates an example of execution of Action $crab(3\pi/4)$, with the corresponding predictions generated by the single and multi-output GPs.

Because the state-of-the-art approaches can only consider uncertainty on one dimension at a time, the other dimension has to be assumed deterministic. For example, when considering heading uncertainty, the distance travelled is assumed to be $\overline{\Delta s}_{dist,a}$ for action a . Similarly, when considering distance uncertainty only, the change in heading to the end position of the rover is assumed to be $\overline{\Delta s}_{head,a}$ for action a . However, whenever the uncertainty is significant in both dimensions, this can generate large prediction errors, as in the figure. This further shows the benefits of using multi-output learning to learn the mobility prediction model.

Table 2: Position errors (m) for single-output and multi-output GP predictions, using the single-line features.

Uncertainty considered:	Single-output GP				Multi-output GP	
	heading		distance		heading & distance	
Action a	mean	std	mean	std	mean	std
$crab(0\pi)$	0.0199	0.0200	0.0148	0.0181	0.0148	0.0181
$crab(\pm\pi/4)$	0.0139	0.0241	0.0108	0.0157	0.0108	0.0157
$crab(\pm\pi/2)$	0.0278	0.0290	0.0132	0.0103	0.0132	0.0103
$crab(\pm 3\pi/4)$	0.0080	0.0142	0.0063	0.0102	0.0034	0.0040
$crab(\pm\pi)$	0.0261	0.0065	0.0185	0.0062	0.0024	0.0021




Fig. 4: A specific example of predictions of the end position of the rover for an execution of Action $crab(3\pi/4)$, compared with the ground truth (s' , shown as a blue solid square). The starting position of the rover is shown as s . All predictions s'_* are indicated by the hollow, black squares. (a) Single output GP prediction considering uncertainty in distance only. (b) Single output GP prediction considering uncertainty in heading only. (c) Multi-output GP prediction.

Multi-line Features Table 3 shows the position errors obtained when predicting the outcomes of each action a using multi-output GPs with the proposed enhanced strategy to collect features that better represent relevant terrain profiles (i.e. *multi-line features*), compared with the predictions made using the state-of-the-art approach with single-output GPs (which considers uncertainty in heading or distance, respectively). The results show that, again, the proposed multi-output learning approach is more accurate than both state-of-the-art single-output techniques.

In addition, comparing the last columns of Tables 2 and 3 shows that the new strategy consisting in capturing features over multiple lines on the terrain rather a single line leads to more accurate and more consistent mobility predictions. This indicates that the multi-line features provide more appropriate information on the terrain profiles than the features computed over a single line.

In summary, the experimental results in this section validate that the learned mobility prediction model is much more accurate when using: 1) multi-output GPs, and 2) features capturing terrain profiles characteristics over multiple lines rather a single line.

Table 3: Position errors (m) for single-output and multi-output GP predictions, using the multi-line features.

Uncertainty considered:	Single-output GP				Multi-output GP	
	heading		distance (m)		heading & distance (m)	
Action a	mean	std	mean	std	mean	std
$crab(0\pi)$	0.0198	0.0201	0.0055	0.0106	0.0054	0.0106
$crab(\pm\pi/4)$	0.0132	0.0228	0.0095	0.0170	0.0062	0.0106
$crab(\pm\pi/2)$	0.0266	0.0307	0.0092	0.0083	0.0073	0.0090
$crab(\pm3\pi/4)$	0.0070	0.0141	0.0041	0.0046	0.0030	0.0034
$crab(\pm\pi)$	0.0064	0.0083	0.0189	0.0058	0.0015	0.0020

Table 4: Summary of all experimental runs

Uncertainty considered	Total runs	Successful runs	Failed runs (%)
None (deterministic)	10	0	10 (100%)
Heading	11	10	1 (9.1%)
Joint (proposed)	10	10	0 (0%)

7 Experimental Results - Planning

In this section, we validate the use of the enhanced SMPM in motion planning and execution, using Mawson on unstructured, partially deformable terrain. We evaluated the performance of planning and execution using the SMPM learned by the approach in this paper, which considers the joint heading and distance uncertainties for crabbing actions, and yaw uncertainty for rotational actions, using multi-output GPs trained with the new features taken over multiple lines on the terrain.


We compared this performance with one of the state-of-the-art methods, which learns the SMPM using a single-output GP, trained with features taken over a single line ahead the rover. In the experiments of this section we use the method that considers heading uncertainty for crabbing actions, and yaw uncertainty for rotational actions. We chose to consider heading uncertainty rather than distance uncertainty for this comparison because the results in [3] indicated that the former had more impact on the performance of planning and execution, both in terms of reliability and cost reduction.

For reference, we also compared with a control method that uses a deterministic mobility prediction model, based on the $\overline{\Delta s_a}$ values from the training data (similarly to the experiments in [3]).


Using a depth sensor, point clouds of the terrain were captured at known locations to generate a DEM, which allowed for the computation of the cost map and λ features (both single-line and multi-lines) representing the terrain profiles. We then generated the SMPMs using the different approaches to be compared. Finally, we defined a common goal area on the map and built policies to reach this goal, using DP with each type of SMPM. Once these policies were obtained, an experimental run corresponded to the rover following the given policy from a starting location s_0 on the map. We executed multiple experimental runs for each type SMPM (i.e. for each policy), to account for the stochastic nature of the process.

Table 4 summarises the amount of runs performed for each method. Each *successful run* corresponds to one full trajectory executed by the rover until it reached the goal. *Failed runs* correspond to cases when the rover failed to reach the goal, because it was stuck on rocks or in a crater (located approximately at $(x, y) = [6, 0]$ in Fig. 5), and/or had its wheels bogged in loose soil. The numbers in parentheses in the table show the percentage of number of failed runs over the total number of runs.

When no uncertainty was considered, all runs failed. This is because the rover tried to traverse into loose soil sections of the crater, where wheels often become stuck. Conversely, when using both approaches that consider uncertainty the rover successfully escaped the crater in most cases by traversing parts of the crater with more rocks in the



(a) State-of-the-art approach (heading uncertainty and single-line features)



(b) Proposed approach (joint uncertainty and multi-line features)

Fig. 5: Multiple executed paths obtained when using: (a) the policy generated from the SMPM learned using the state-of-the-art approach, and (b) the policy generated from the enhanced SMPM, using the proposed approach. The background shows the cost map, coloured by cost value (see the colour bar on the top right).

Table 5: Statistics for successful experimental runs

Uncertainty Considered:	Mean length (m)	No. runs	Cost			No. actions	
			Mean	Std	Reduction	Mean	Std
Heading uncertainty	15.87	10	0.2984	0.0935	N/A	67.20	7.66
Joint uncertainty (proposed method)	16.04	10	0.1810	0.0196	39.34%	46.50	5.323

ground. Note that for each method that considers uncertainty, we used the same number of successful runs (10) in order to generate comparable statistics.

Figure 5 illustrates a subset of the executed paths, shown over the cost map used for these experiments. For clarity, we only show five executed paths for each successful method, chosen randomly. We can observe that the executed paths are fairly spatially consistent. It appears that there is less spatial variation between the multiple executed paths when only heading uncertainty was considered, compared with the proposed approach where joint uncertainty was considered.

The statistics of actual total cost integrated over each executed path, path length, and number of actions executed for the successful experimental runs are shown in Table 5. It can be observed that the mean and standard deviation of the total cost accumulated along the path executions are significantly lower with the proposed approach (39.34% reduction). This cost reduction can be considered as highly statistically significant, since performing the significance test from [20] gave a p -value of 0.00005.

The average lengths of executed paths are comparable (15.87m and 16.04m). However, the rover had to execute a much smaller number of actions when using the proposed approach. This indicates that the action executions were much more efficient in average. Furthermore, the reduced standard deviation on the number of executed actions suggests that the action executions were more consistent.

Overall, the experimental results show that the rover greatly benefits from using SMPMs generated by the proposed approach when planning and executing policies, especially in terms of cost accumulated over the executed paths, and in terms of action efficiency in practice.

8 Conclusion and Future Work

We have presented a new method for mobility prediction based on multi-output GP regression. We evaluated our method experimentally in comparison with our previous single-output GP method and also a no-uncertainty control condition.

Our experiments show that mobility prediction with multi-output GPs is clearly beneficial for navigation tasks. In the control condition the rover failed to reach its goal in all trials, whereas no trials failed in the multi-output condition. The multi-output condition resulted in better path execution as measured by fewer total actions.

These results further validate the role of mobility prediction in achieving safe, reliable navigation for planetary rovers. Important avenues for future work include the consideration of additional sources of uncertainty, such as localisation.

Acknowledgments

This work was supported in part by the Australian Centre for Field Robotics (ACFR) and the NSW State Government.

References

1. Johnson, J., ed.: Mars Scientific Goals, Objectives, Investigations, and Priorities. Mars Exploration Program Analysis Group (MEPAG), NASA/JPL (2010) <http://mepag.jpl.nasa.gov/reports/index.html>.
2. McAllister, R., Peynot, T., Fitch, R., Sukkariéh, S.: Motion planning and stochastic control with experimental validation on a planetary rover. In: Proc. of IEEE/RSJ IROS. (2012)
3. Peynot, T., Lui, S.T., McAllister, R., Fitch, R., Sukkariéh, S.: Learned stochastic mobility prediction for planning with control uncertainty on unstructured terrain. Journal of Field Robotics (2014) In press.
4. Ishigami, G., Kewlani, G., Iagnemma, K.: Statistical mobility prediction for planetary surface exploration rovers in uncertain terrain. In: Proc. of IEEE ICRA. (2010)
5. Krebs, A., Pradalier, C., Siegwart, R.: Adaptive rover behavior based on online empirical evaluation: Rover-terrain interaction and near-to-far learning. Journal of Field Robotics 27(2) (2010) 158–180
6. Brooks, C.A., Iagnemma, K.: Self-supervised terrain classification for planetary surface exploration rovers. Journal of Field Robotics 29(3) (2012) 445–468

7. Papadakis, P.: Terrain traversability analysis methods for unmanned ground vehicles: A survey. *Engineering Applications of Artificial Intelligence* **26**(4) (2013) 1373–1385
8. Goldberg, S.B., Maimone, M.W., Matthies, L.: Stereo vision and rover navigation software for planetary exploration. In: *Proc. of IEEE Aerospace Conference*. (2002)
9. Helmick, D.M., Cheng, Y., Clouse, D.S., Matthies, L.H., Roumeliotis, S.I.: Path following using visual odometry for a mars rover in high-slip environments. In: *Proc. of IEEE Aerospace Conference*. (2004)
10. Chhaniyara, S., Brunskill, C., Yeomans, B., Matthews, M., Saaj, C., Ransom, S., Richter, L.: Terrain trafficability analysis and soil mechanical property identification for planetary rovers: A survey. *Journal of Terramechanics* **49**(2) (2012) 115–128
11. Iagnemma, K., Kang, S., Shibly, H., Dubowsky, S.: Online terrain parameter estimation for wheeled mobile robots with application to planetary rovers. *IEEE Transactions on Robotics* **20**(5) (2004) 921–927
12. Helmick, D., Angelova, A., Matthies, L.: Terrain adaptive navigation for planetary rovers. *Journal of Field Robotics* **26**(4) (2009) 391–410
13. Karumanchi, S., Iagnemma, K.: Reactive control in environments with hard and soft hazards. In: *Proc. of IEEE/RSJ IROS*. (2012)
14. Rasmussen, C.E., Williams, C.: *Gaussian Processes for Machine Learning*. MIT Press (2006)
15. Boyle, P., Frean, M.: Dependent gaussian processes. In Saul, L., Weiss, Y., Boutou, L., eds.: *Advances in Neural Information Processing Systems (NIPS)*. Volume 17. MIT Press (2005)
16. Bonilla, E., Chai, K.M., Williams, C.: Multi-task gaussian process prediction. In Platt, J., Koller, D., Singer, Y., Roweis, S., eds.: *Advances in Neural Information Processing Systems (NIPS)*. Volume 20. MIT Press (2008)
17. Alvarez, M.A., Lawrence, N.D.: Sparse convolved gaussian processes for multi-output regression. In Koller, D., Schuurmans, D., Bengio, Y., Bottou, L., eds.: *Advances in Neural Information Processing Systems (NIPS)*. Volume 21. MIT Press, Cambridge, MA (2009) 57–64
18. LaValle, S.: *Planning Algorithms*. Cambridge University Press (2006)
19. Tarokh, M., McDermott, G.J.: Kinematics modeling and analyses of articulated rovers. *IEEE Transactions on Robotics* **21**(4) (2005) 539–553
20. Moore, D.S., McCabe, G.P.: *Introduction to the Practice of Statistics*. 2nd edn. Freeman and co., New York (1993)