ACFR
AUSTRALIAN CENTRE FOR FIELD ROBOTICS

Experimental Learning for Traversability Estimation and Stochastic Motion Planning on a Planetary Rover

Thierry Peynot, Ken Ho, Angela Lui, Rowan McAllister, Robert Fitch and Salah Sukkarieh

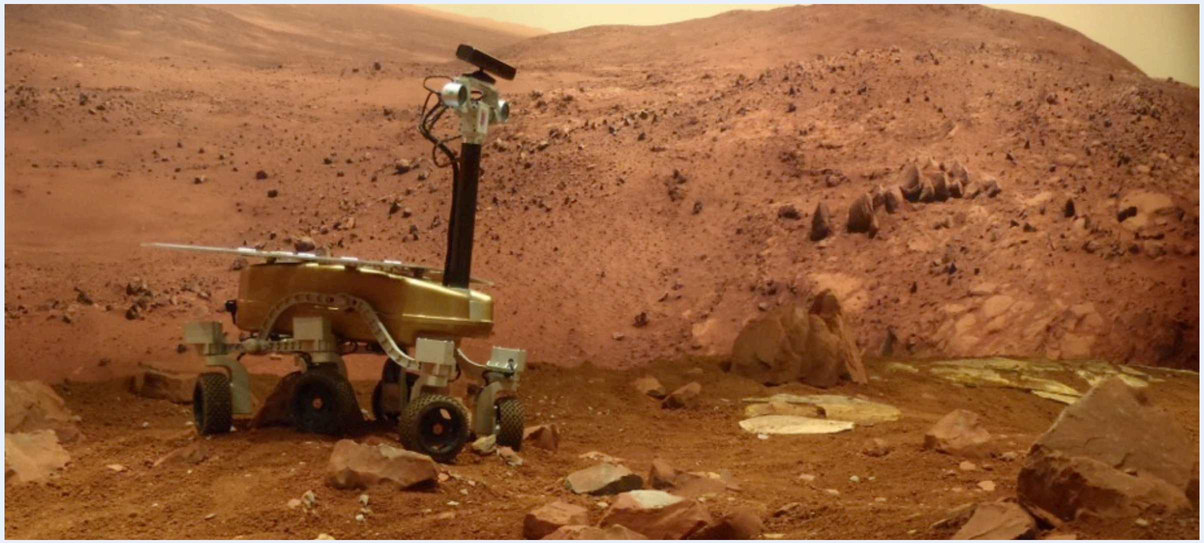

Program

Uncertainty

- Localisation
- Perception (incomplete & uncertain maps) ➔ Traversability
- Control (uncertain outcomes) ➔ Planning

Anticipate (impact of) terrain deformation

Outline:

1. Traversability estimation from incomplete exteroceptive sensing data via experimental learning
2. Motion planning and stochastic control with mobility prediction model learnt from experience

ACFR
AUSTRALIAN CENTRE FOR FIELD ROBOTICS

Traversability Estimation for a Planetary Rover via Experimental Kernel Learning in a Gaussian Process Framework

Ken Ho
Thierry Peynot
Salah Sukkarieh

[ICRA 2013]

Terrain traversability estimation

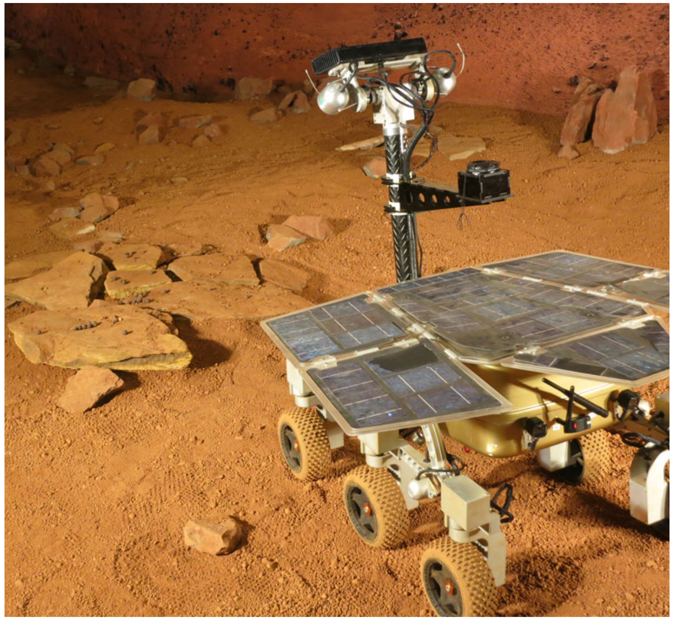

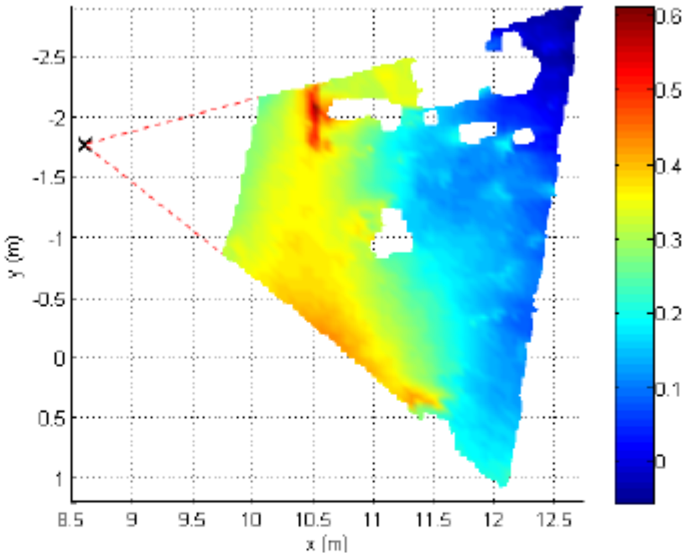

- Objective: Estimate attitude and configuration for a rover
  - Indication of traversability
- Heterogeneous terrain
- Occlusion/missing data in perception
  - Field of view/sensor placement

Previous work

DEM-Kin

- Kinematic modeling on geometric terrain model (Lacroix et. al.)
  - Dependent on geometry of vehicle and terrain
  - All 6 wheels contact ground
    - Need data on all 6 wheels

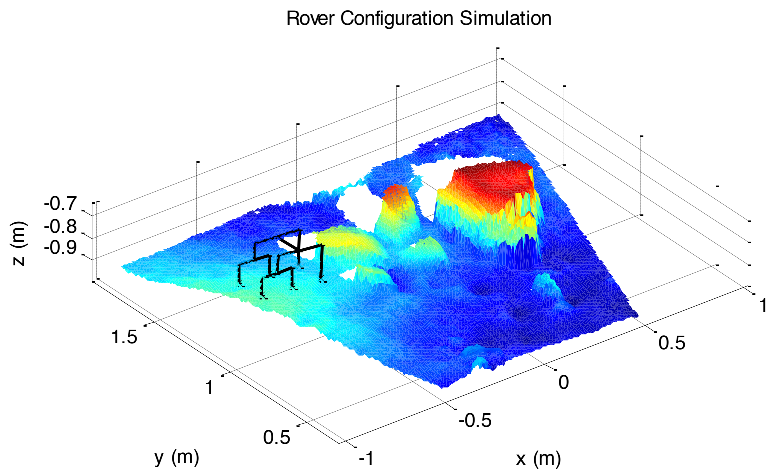

Previous work (cont.)

- Classifying terrain types and predicting slip based on terramechanic model (Iagnemma et. al., Helmick et. al.)
- Improving geometric terrain model (Vasudevan et. al.)
  - Using GP estimates in areas with little/no data

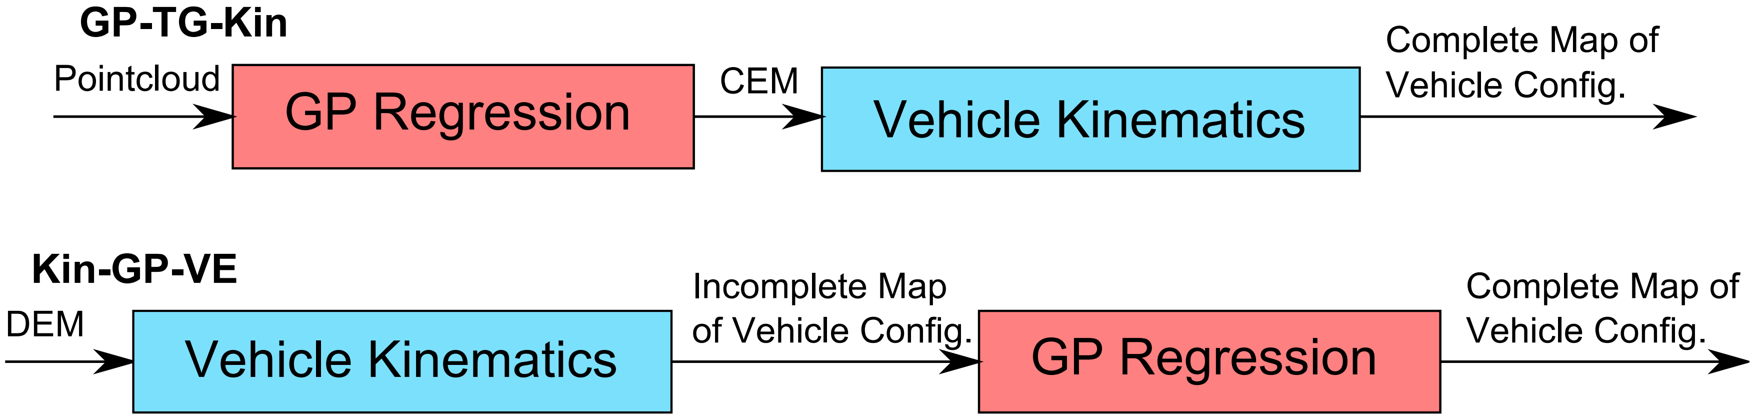

→ This work generates a complete representation of terrain traversability with uncertainty

Proposed Framework – Kin-GPVE

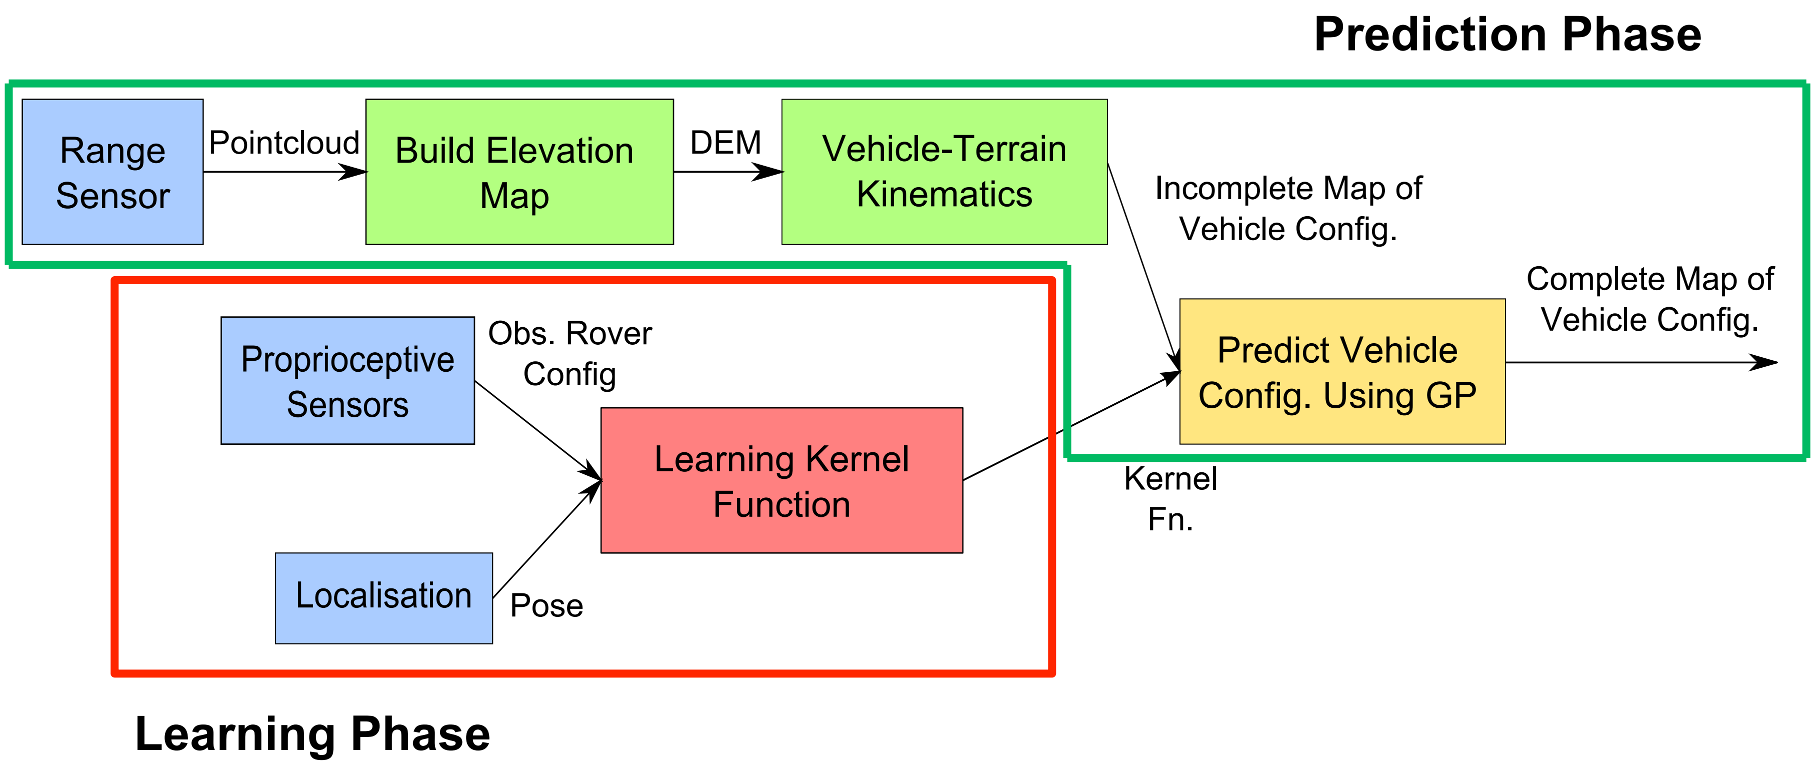

Motivation for Learning Covariance Function

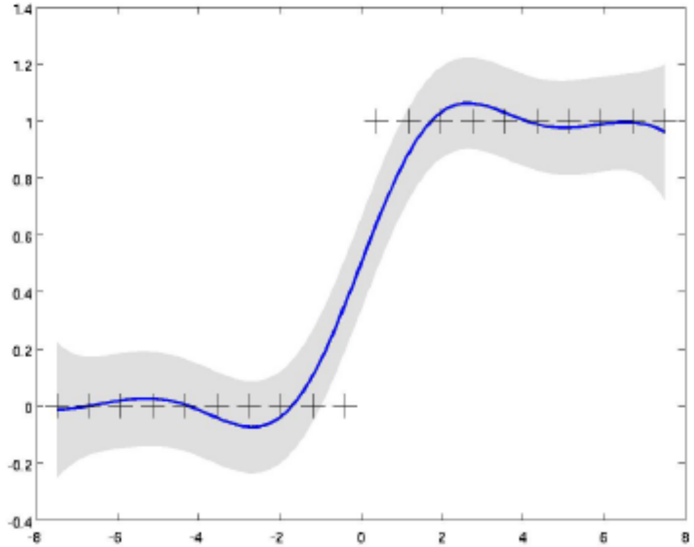

(a) GP modeling of data using SQEXP kernel

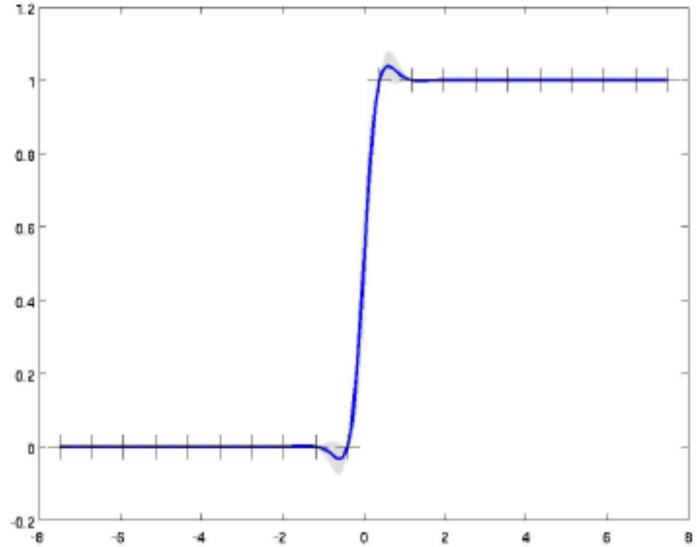

(b) GP modeling of data using NN kernel

- Learning a covariance function suitable to problem at hand
  - Learn from real data
  - Better representation of vehicle configuration evolution
  - Propagation of uncertainty
- Benefit: Explicitly considering vehicle state evolution and propagation of uncertainty in learning framework

(Image source: Vasudevan et. al., JFR 2009)

Overview of Learning Process

1. Estimate the covariance matrix using Regularized Expectation Maximization
2. Combining covariance matrices using Maximum Entropy Covariance Selection (MECS)
3. Generalise the covariance matrix into a function for use in GP framework

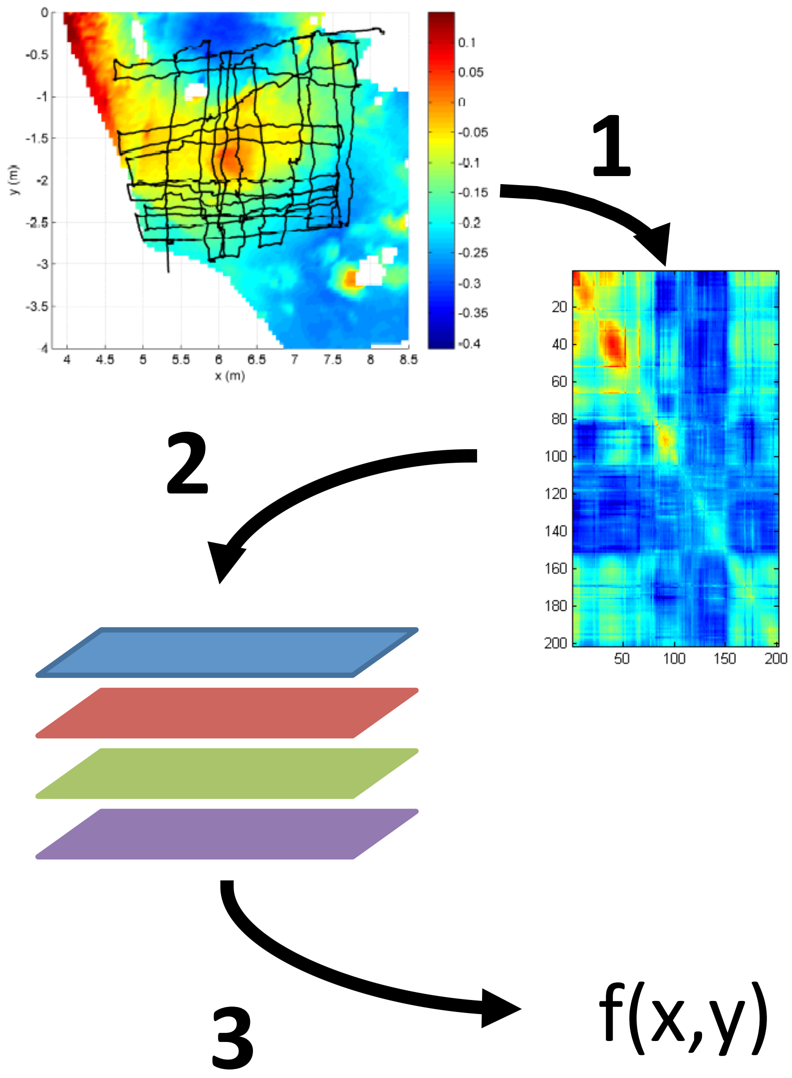

Experimental Setup

- Marsyard at Powerhouse Museum, Sydney
  - Mars Analogue Terrain
- Platform
  - Rocker-bogie chassis
  - Sensors:
    - Kinect RGB-D camera
    - IMU
    - Hall effect encoders
    - Intersense IS-1200 tracking system
- Experiments conducted at slow speed

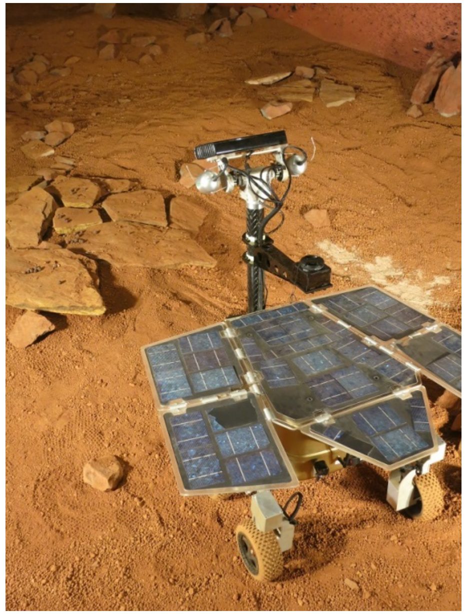

Validation Strategy

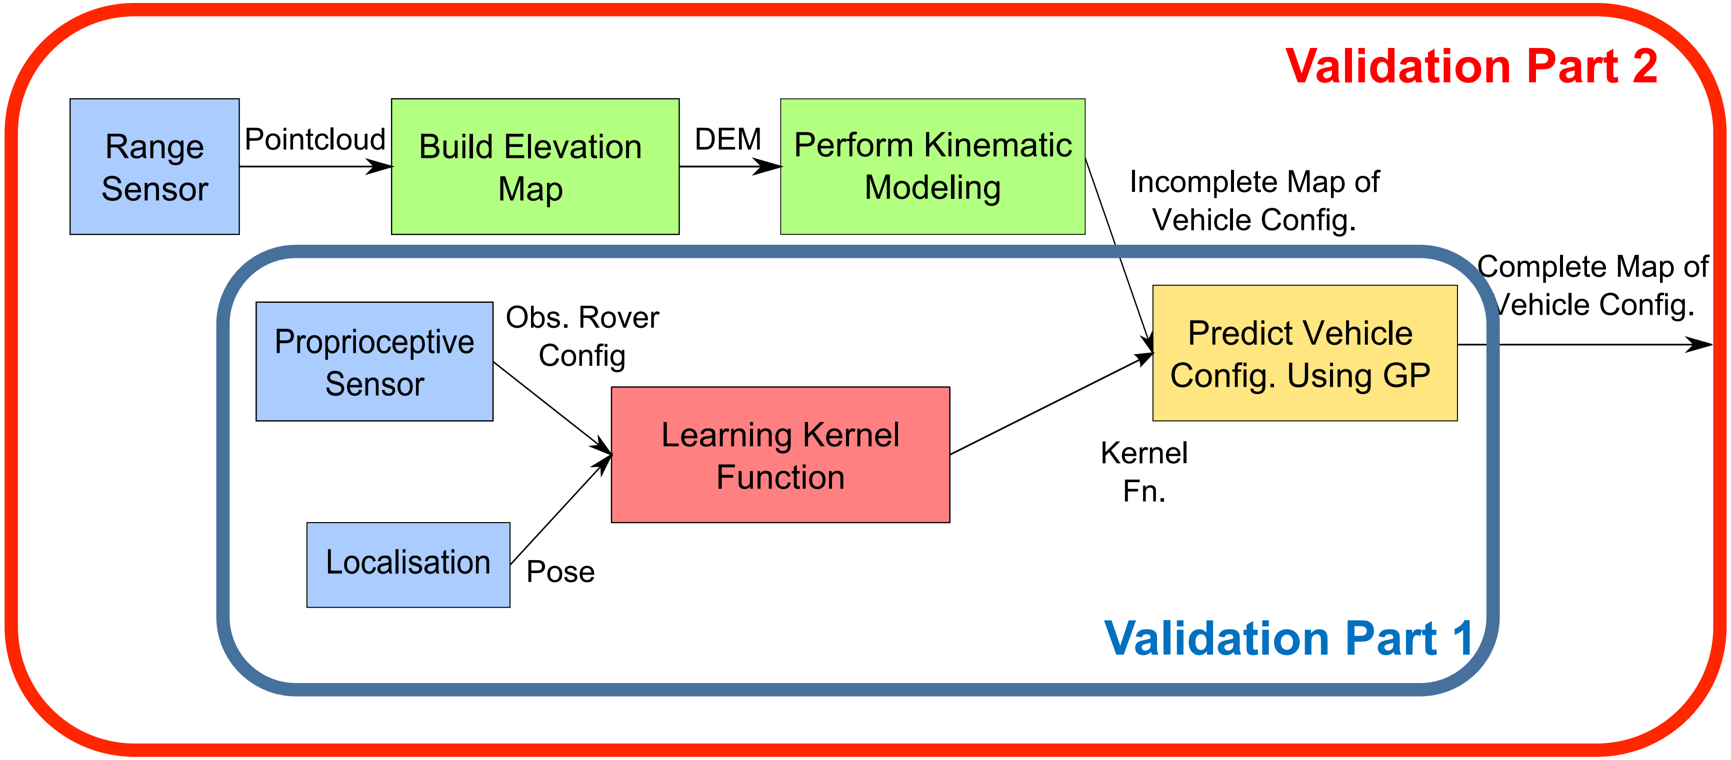

Validation Part 1: Estimating Vehicle Configuration Using Proprioceptive Data

- Performance of learnt kernel function vs. state-of-the-art kernel functions

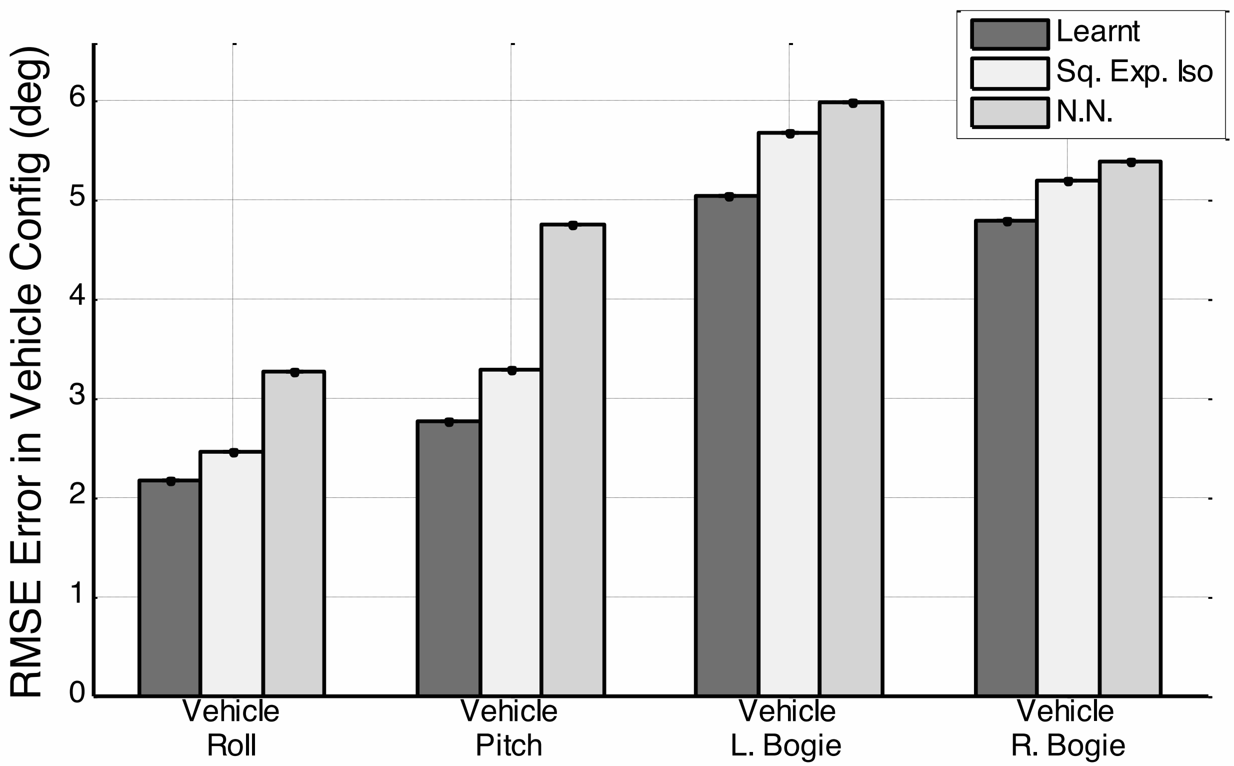

Validation Part 2: Estimating Vehicle Configuration Using Exteroceptive Data

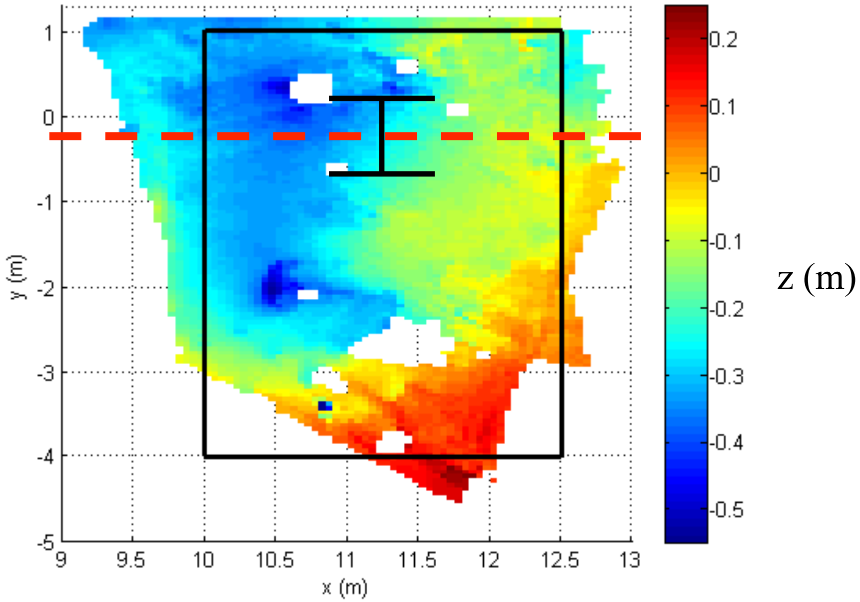

- Incomplete DEM from exteroceptive data
  - Areas with no data are white
  - Occlusions in elevation map ≠ Occlusions in vehicle config.
    - At least 1 wheel touches the ground with no data

Validation Part 2 (cont.)

- Prediction comparison of different methods

terrain deformation

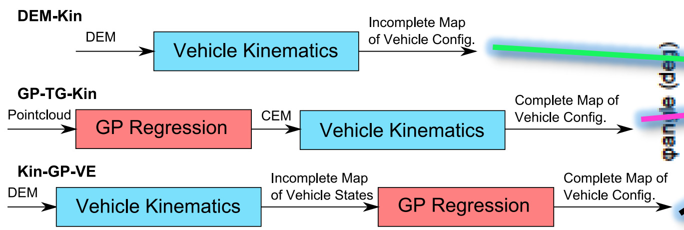

roll

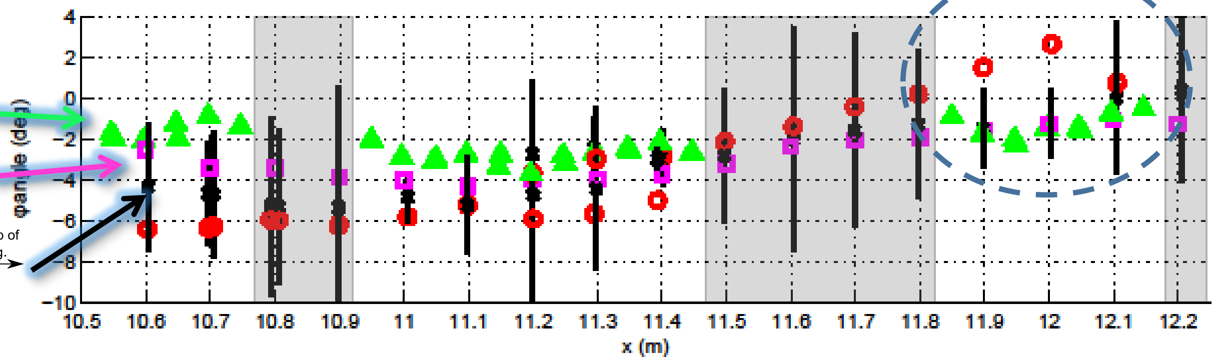

pitch

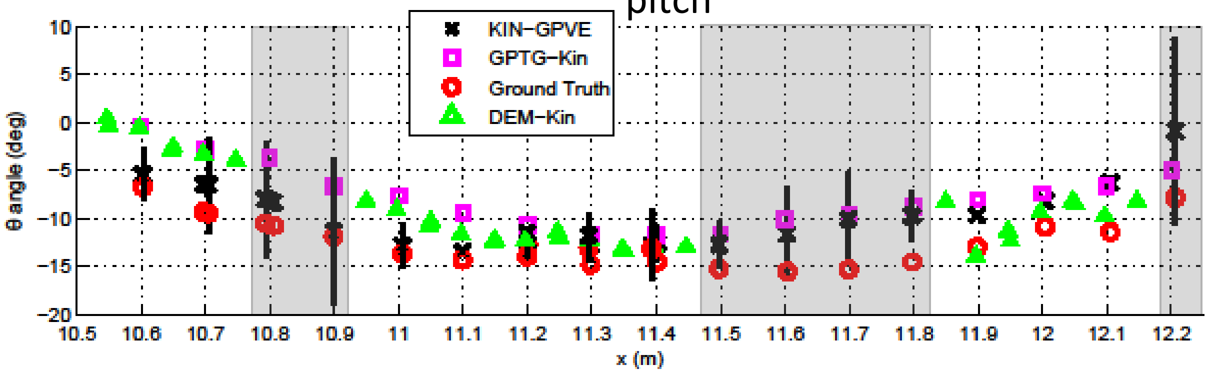

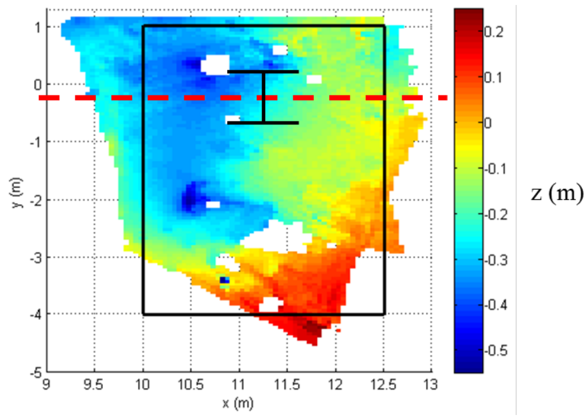

Kin-GP-VE Summary

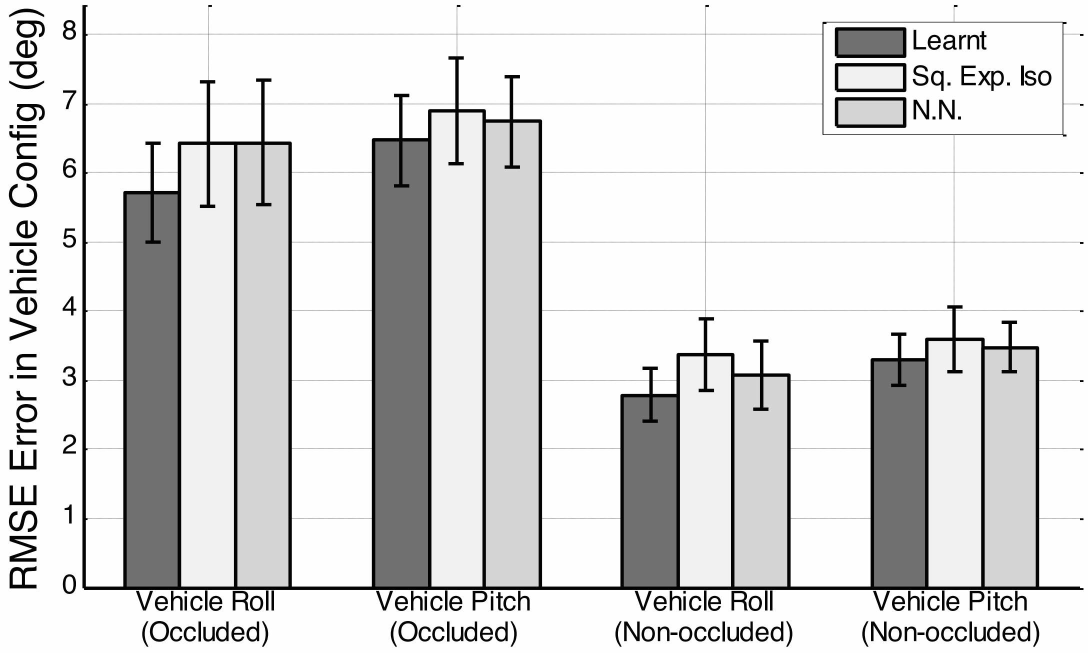

- Consistent improvements in estimating vehicle configuration over state-of-the-art kernel functions
  - Learning kernel function most suitable for problem at hand
- Still restricted by simplifications and assumptions in kinematic model

Deformable Terrain…

- Need method for predicting vehicle configuration angles of the rover on deformable terrain
  - Refine estimate from Kin-GP-VE to account for terrain deformation
  - Include dynamic influences towards vehicle configuration

ACFR
AUSTRALIAN CENTRE
FOR FIELD ROBOTICS

A Near-to-Far Non-Parametric Leaning Approach for Estimating Traversability in Deformable Terrain

Ken Ho, Thierry Peynot, Salah Sukkarieh

[Submiitted to IROS 2013]

Different Approaches

- Rigid Terrain Traversability Estimation (R-TTE)

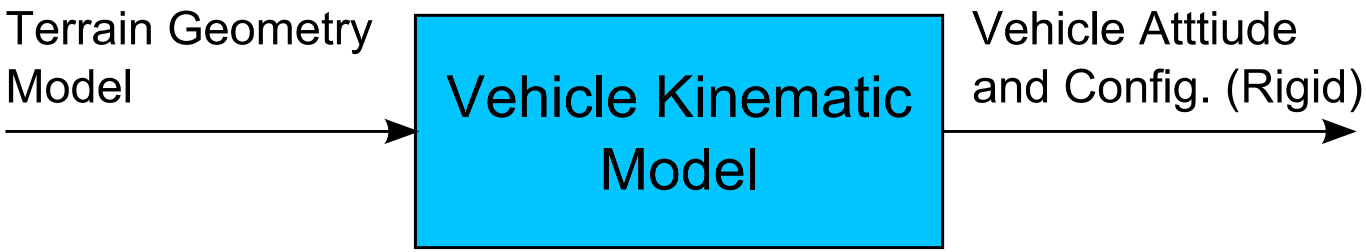

- Rigid to Deformable Terrain Traversability Estimation (R2D-TTE)

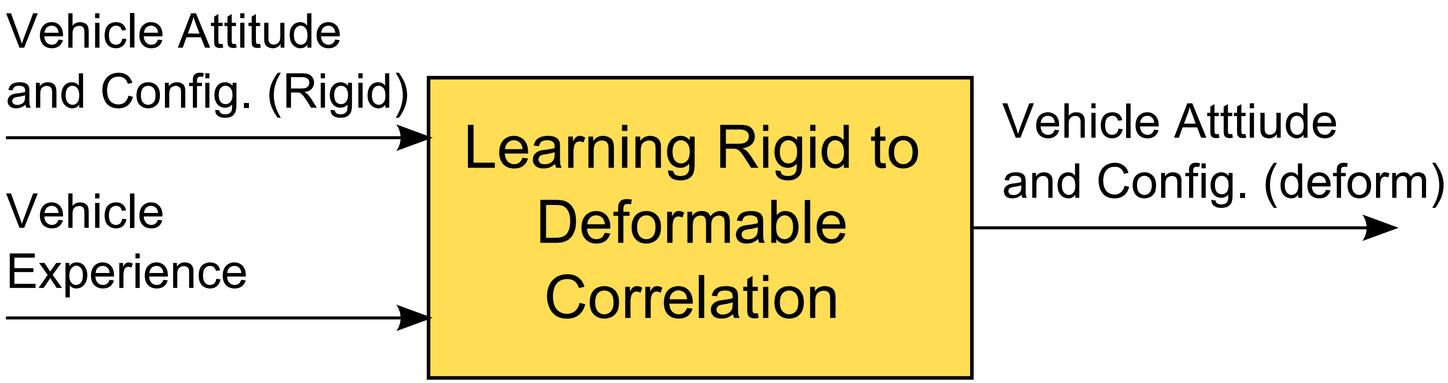

Learning correlations in R2D-TTE

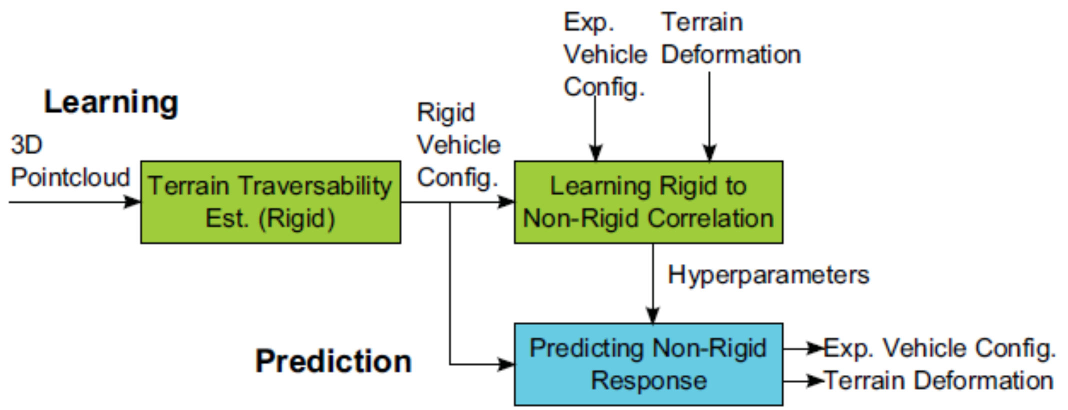

Fig. 3. System Architecture for Rigid to Deformable Terrain Traversability Estimation (R2D-TTE)

1. Predicting vehicle configuration on rigid terrain (R-TTE)
2. Learning correlations between prediction made in (1) and vehicle experience, which include vehicle configuration and terrain deformation from experience

Learning correlations between R-TTE and vehicle experience

- Multiple input GP regression by automatic relevance determination (ARD)
  - Sq-exp, separate length-scale for each input

$$k\left(X, X'\right) = \sigma_f^2 \exp\left(-\frac{1}{2}\left(X - X'\right)^T M \left(X - X'\right)\right) + \sigma_n^2 \delta_{pq} \qquad M = diag(l)^{-2}$$

- Multi-task GP regression
  - Heteroscedastic noise

$$k_q\left(X - z\right) = \frac{S_q |M_q|^{1/2}}{(2\pi)^{p/2}} \exp\left[-\frac{1}{2}\left(X - z\right)^T M_q \left(X - z\right)\right] \tag{9}$$

  - Between 2 outputs

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

$$cov\left[f_q(X), u_r(z)\right] = \int_{-\infty}^{\infty} k_{qr}\left(X - z'\right) k_{u_r u_r}\left(z', z\right) dz'$$

Validation Results

- Predicting vehicle roll over 500 validation points

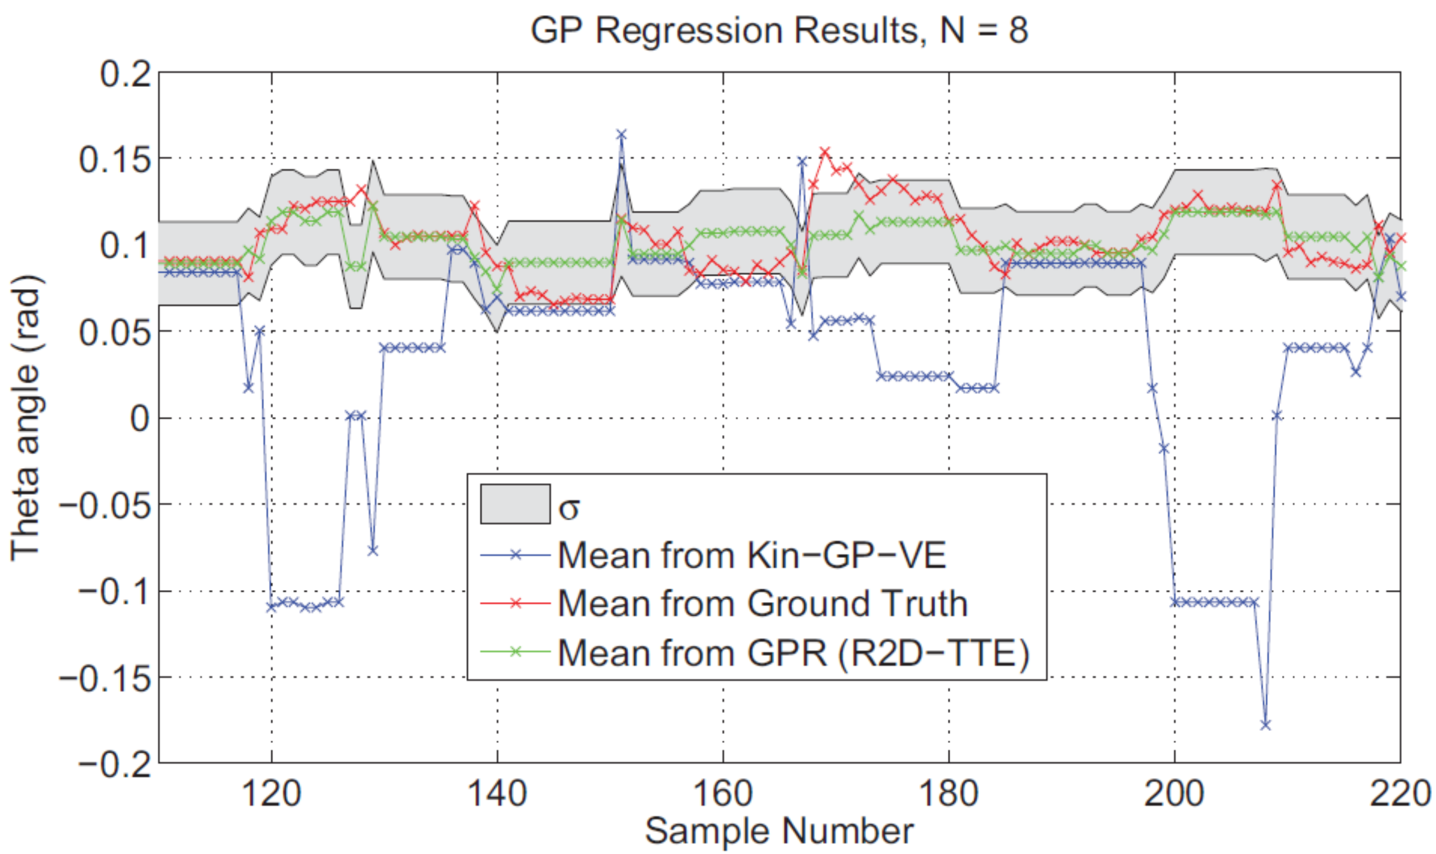

Validation Results (cont.)

- Predicting deformation over 500 validation points

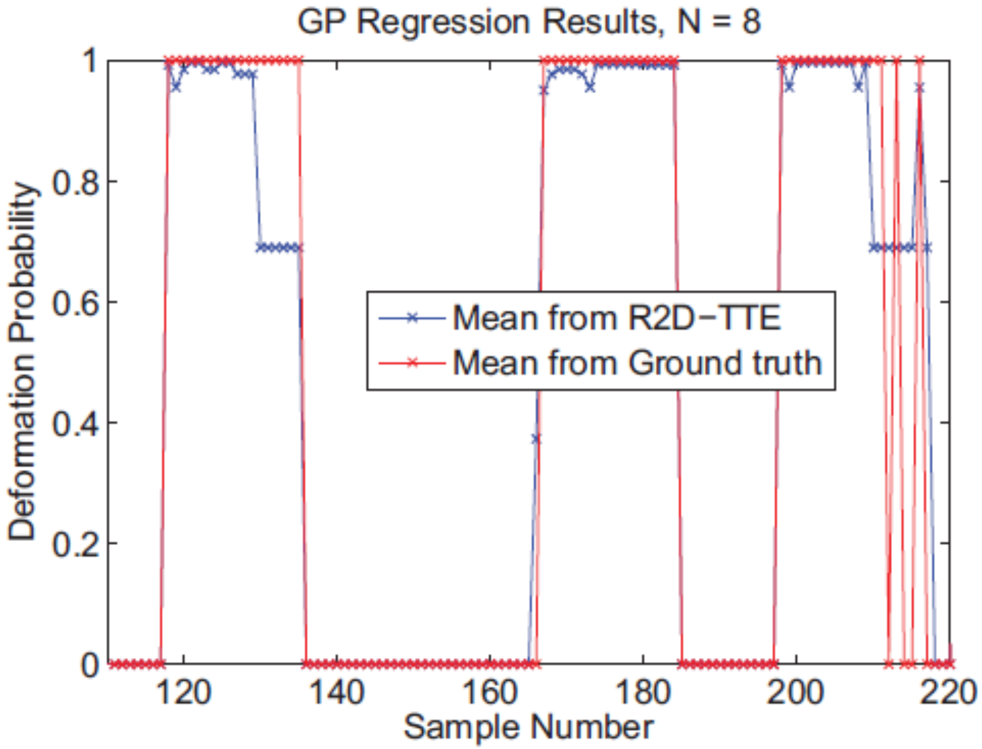

Fig. 7. GP Regression results for predicting $\mathcal{T}_{deform}$ over 500 validation points, zoomed over sample number 110 to 220.

Validation Results (cont.)

- Predicting vehicle roll and left bogie angle over areas with higher deformation than that experienced during training

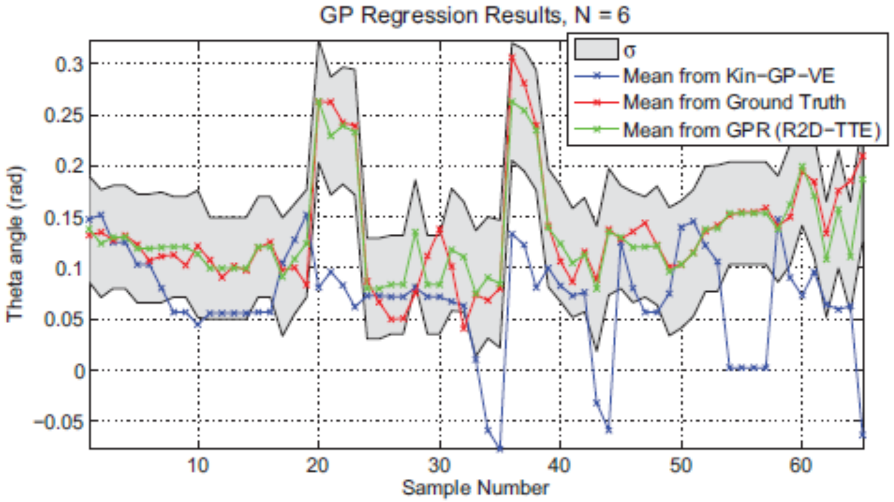

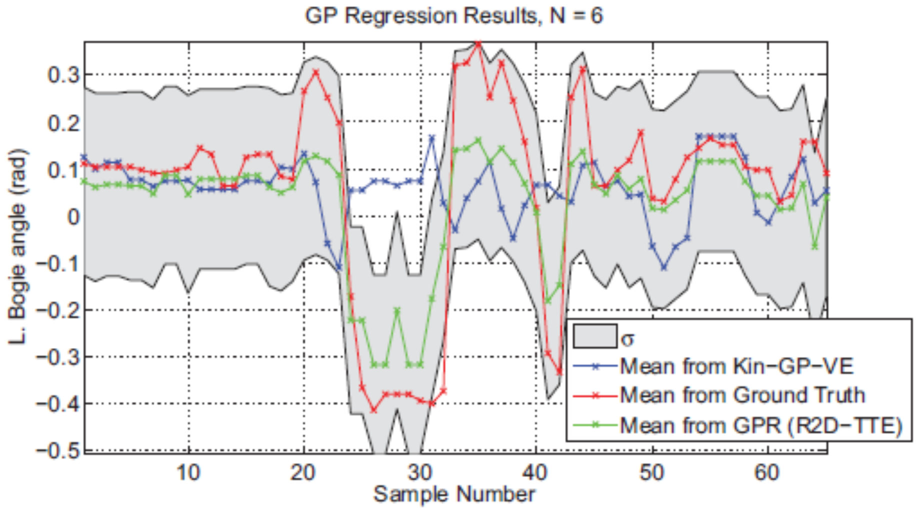

Validation Results (cont.)

- Overall RMSE from experiments

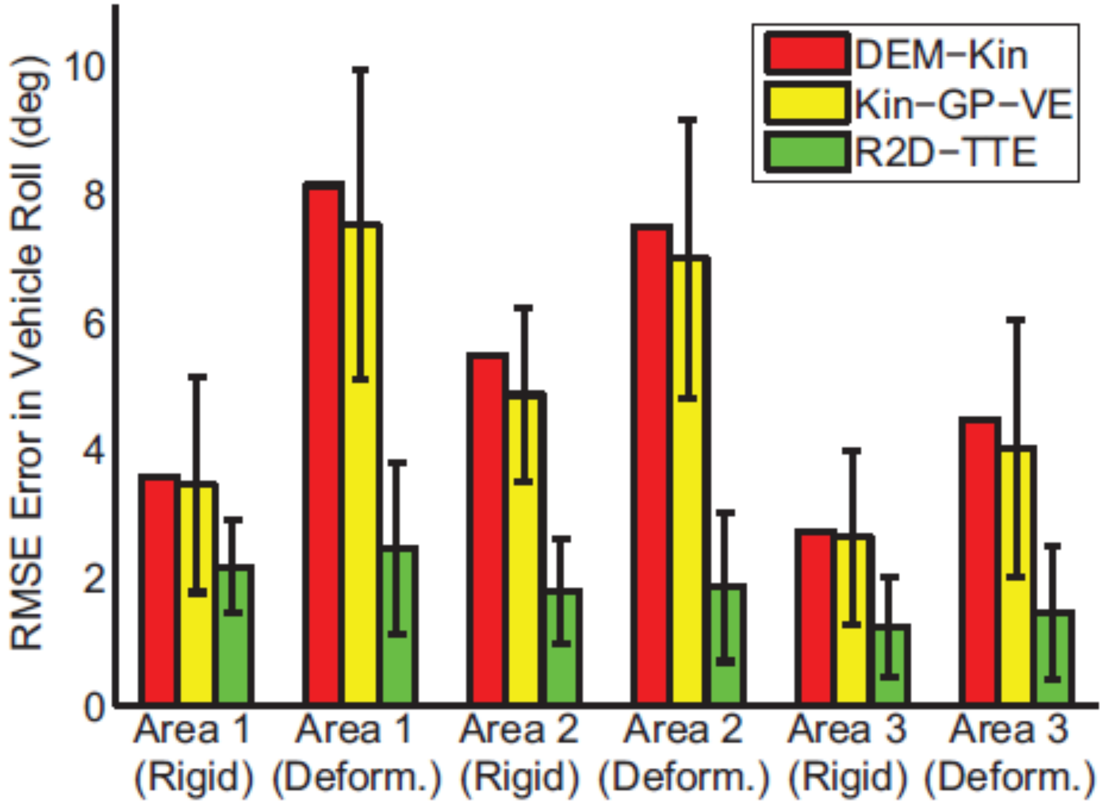

Conclusions and Future Work

- Novel method for predicting vehicle configuration angles of a planetary rover over deformable terrain.
- Will consider terrain descriptors other than geometry that would contribute towards discerning deformable terrain, such as color and texture.
- Require more accurate measurements of the changes in terrain geometry as the rover traverses over it.
  - Can be obtained from an external observation setup, such as a geo-referenced LIDAR or a multi-camera system.

ACFR
AUSTRALIAN CENTRE FOR FIELD ROBOTICS

Motion Planning and Stochastic Control with Experimental Validation on a Planetary Rover

Angela Lui, Rowan McAllister, Thierry Peynot, Robert Fitch and Salah Sukkarieh

[IROS 2012]
[JFR ?]

Motion Planning in Unstructured Terrain

Objective:
to traverse unstructured terrain towards a goal **reliably** and **safely**

Motion planning needs to takes into account uncertainty

- localisation
- mapping
- **control**

Our approach considers uncertainty in control and is practical for real platform on realistic unstructured terrain

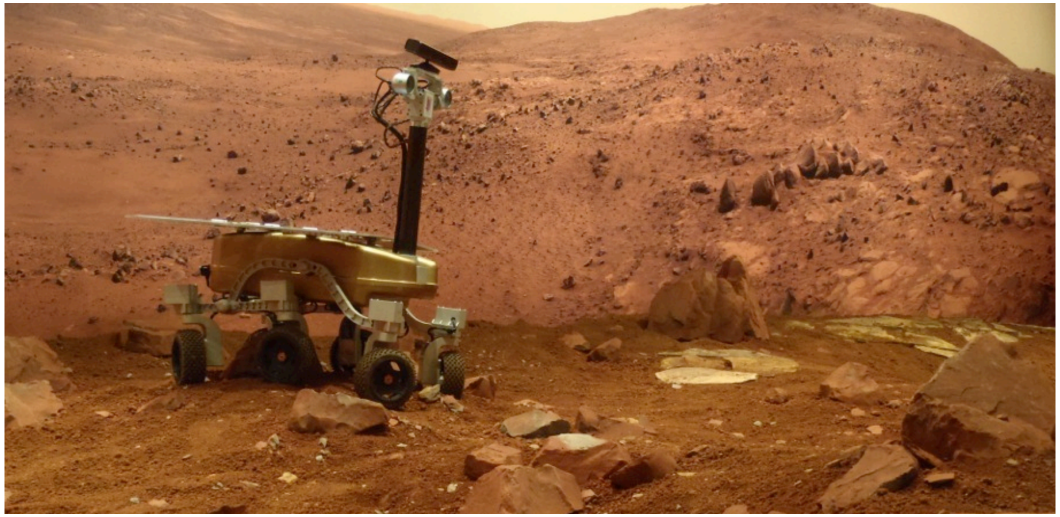

Related Work

Most approaches use a deterministic motion planner (A*, RRT, PRM) to compute candidate paths, assess control uncertainty along each one, execute the "best" candidate.

Using this method, control uncertainly modelled by either:

- expected feedback controller's deviations,

  [Berb 2010][Bry 2011][Greytak 2009][Ishigami 2007][Platt 2010][Patil 2011]
- terramechanics,

  [Ishigami 2010]

Limitations:

- Planning and control are decoupled.
- LQG: assumes homogeneity of control uncertainty.
- Terramechanics: difficult to model non-homogeneous terrain locally

Stochastic Mobility Prediction Model

- Uncertain outcomes of control action executions: deviations in yaw, heading and distance
- Unstructured terrain: given one action, outcomes depend on profile of terrain traversed
- Learn a model of control uncertainty by experience (learn stochastic mobility prediction)

Implementation:

- Learning: Train Gaussian Processes (GP) with multiple action executions and observed terrain profiles during traversals.
- Use GP regression to obtain a stochastic transition function in motion planning (ability to query a predictive distribution of outcomes for any action and terrain profile) $p(\Delta s|\boldsymbol{\lambda}(s,a),a)$
- Prediction: use DEM and kinematics model to predict pitch and roll evolution that will be mapped to action outcomes distribution.

Planning Algorithm

- Stochastic model of mobility prediction (stochastic transition model learnt from experience)
- Compute policies using Dynamic Programming (DP)

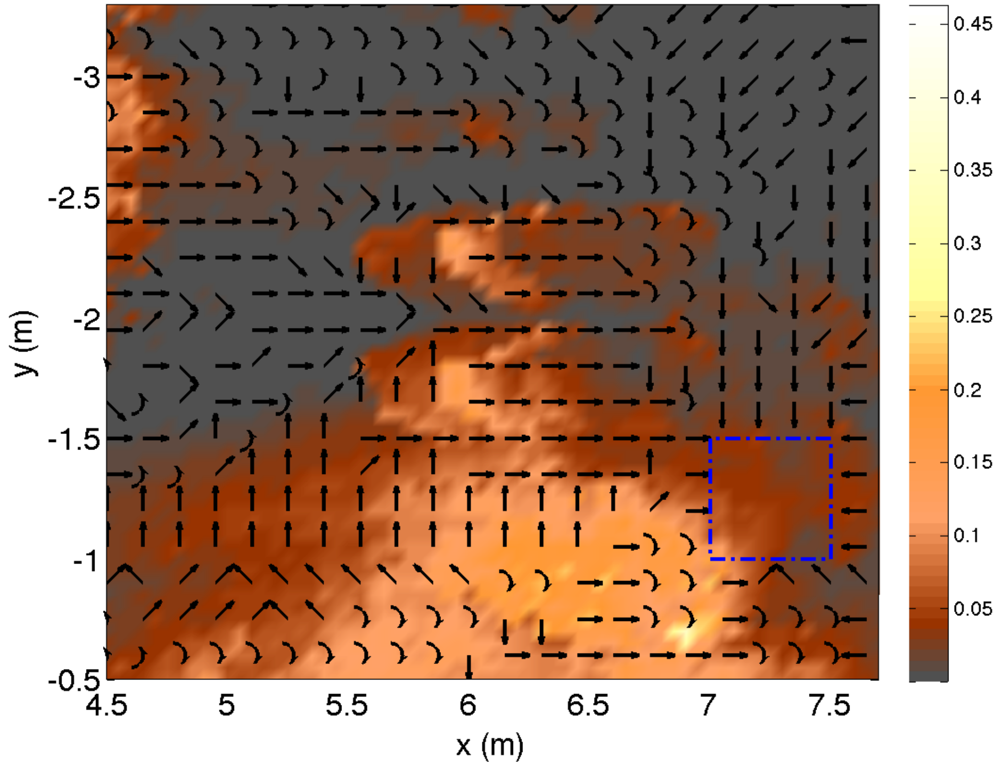

Bellman Eq. for optimal policy:

$$\pi^*(s) = arg\max_a \left\{ \sum_{s'} \underbrace{P(s'|s,a)}_{\text{Transition Function}} \left(\underbrace{R(s',s,a)}_{\text{Reward Function}} + \gamma \underbrace{V(s')}_{\text{Value Function}} \right) \right\}$$

Implementation

- Action set:
  - CRAB actions (8 directions)
  - ROTATE actions (Clockwise, Anti-Clockwise)

- Digital Elevation Map (DEM)
- Method to predict traversability (rover configuration angles, e.g. Kinematic Modelling)
- Cost map: penaliser per action + terrain cost (difficulty)

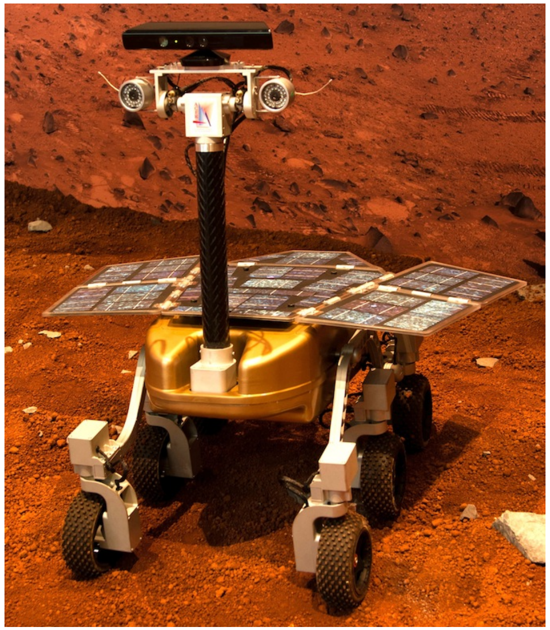

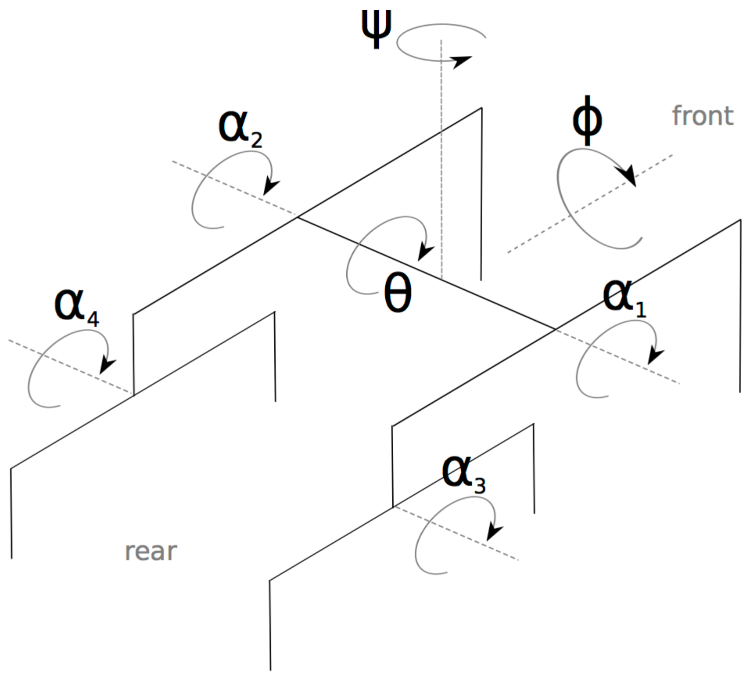

Framework

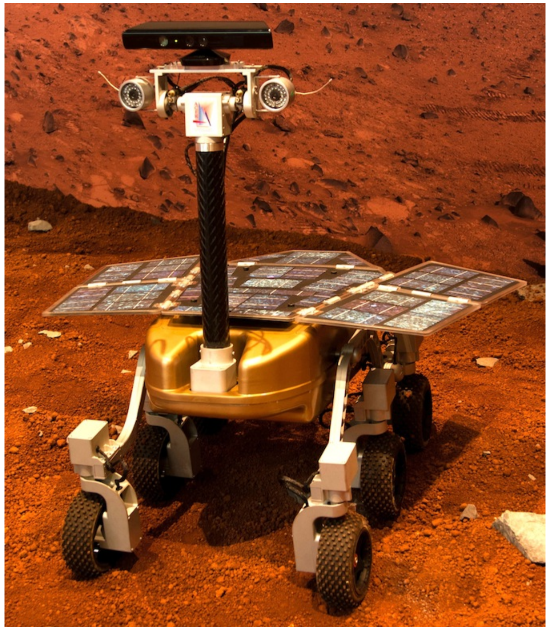

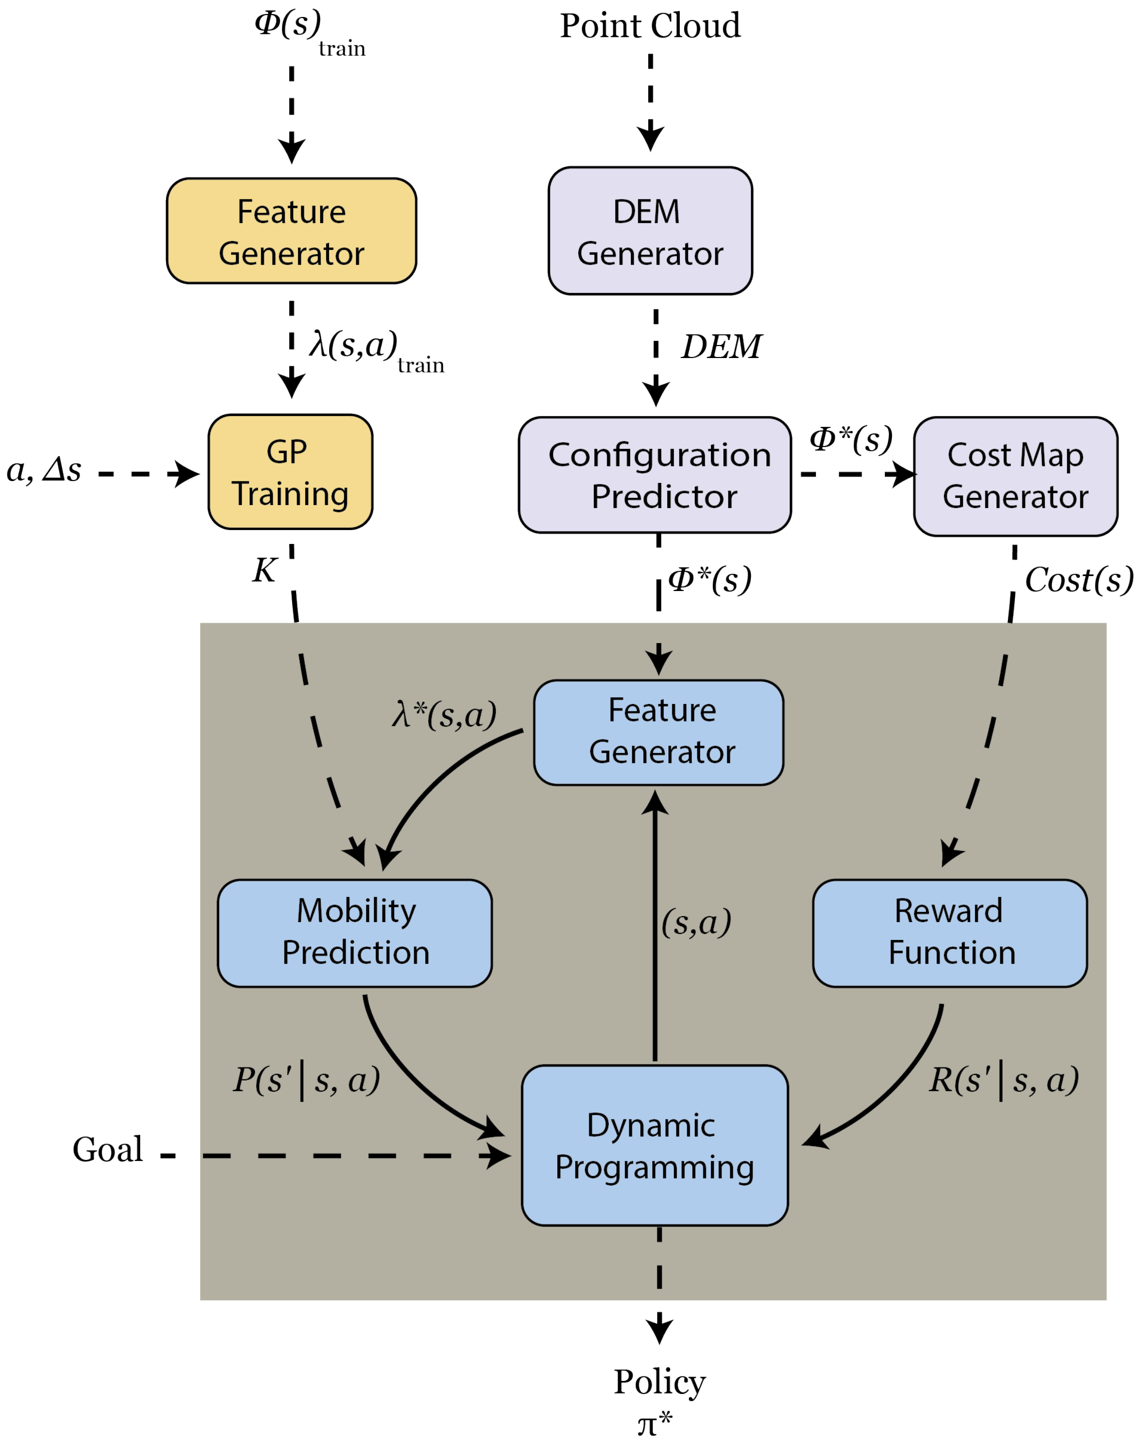

$\Phi(s)_{train}$
Point Cloud
Feature Generator
DEM Generator
$\lambda(s,a)_{train}$
DEM
$a, \Delta s$
GP Training
Configuration Predictor
$\Phi^*(s)$
Cost Map Generator
K
$\Phi^*(s)$
Cost(s)
$\lambda^*(s,a)$
Feature Generator
Mobility Prediction
(s,a)
Reward Function
$P(s' | s, a)$
$R(s' | s, a)$
Goal
Dynamic Programming
Policy
π^*

Learning Mobility Prediction from Proprioception (LfP)

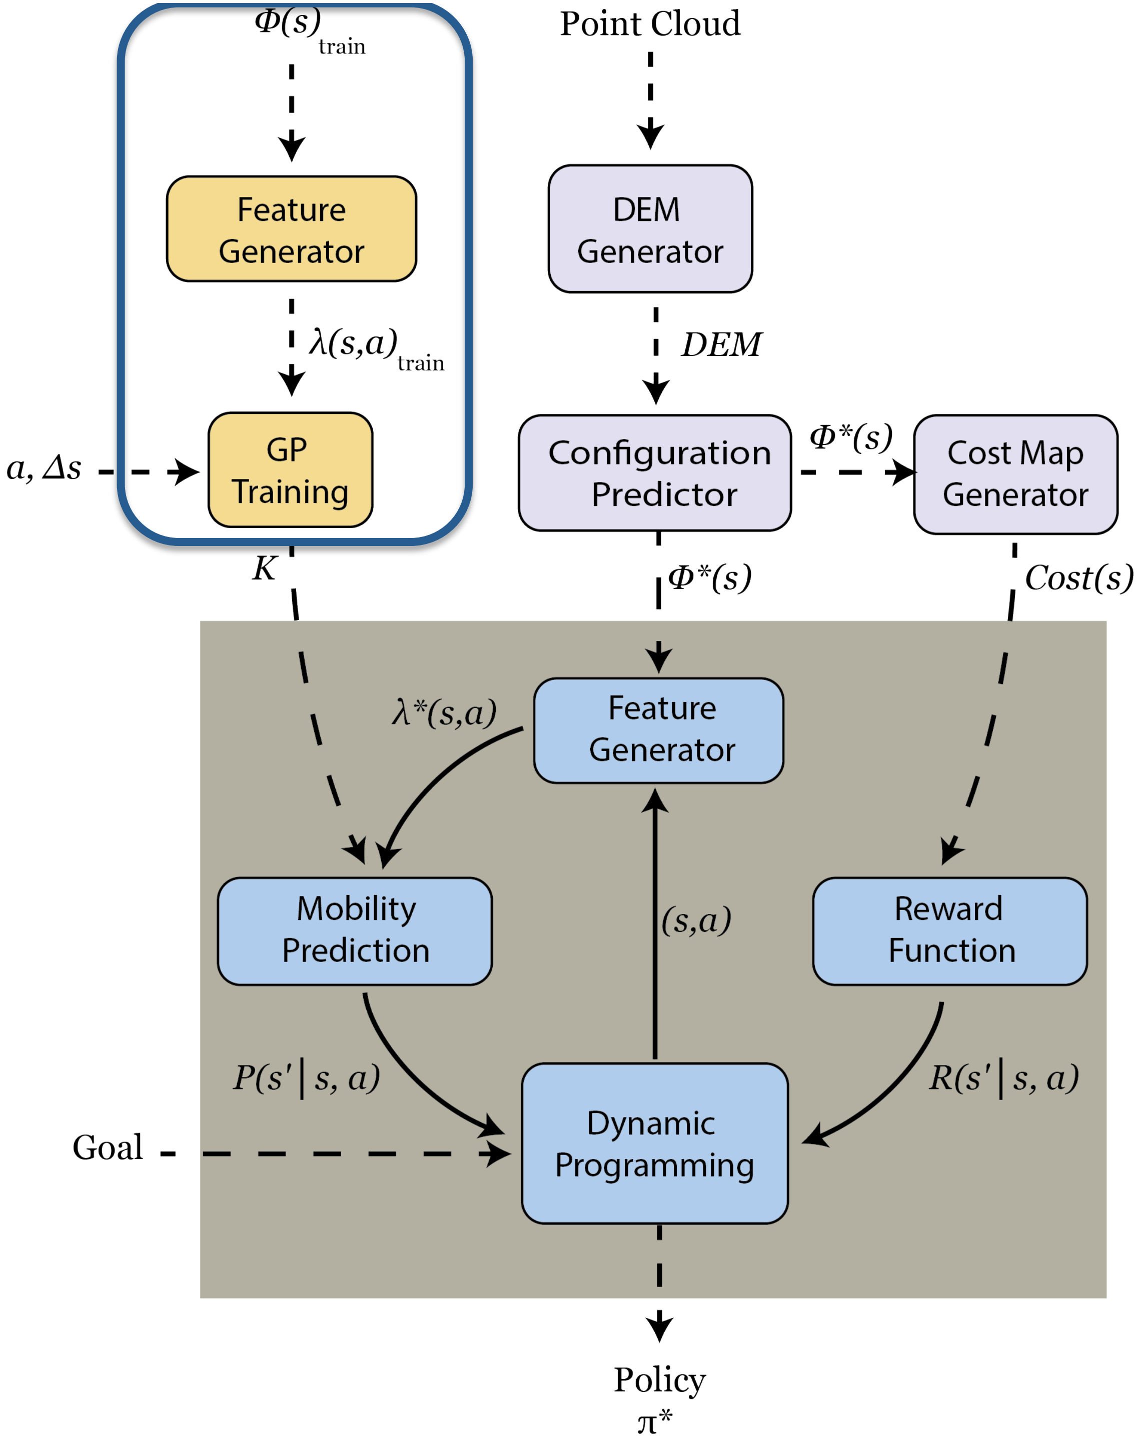

ACFR
AUSTRALIAN CENTRE
FOR FIELD ROBOTICS

Learning Mobility Prediction from Proprioception (LfP)

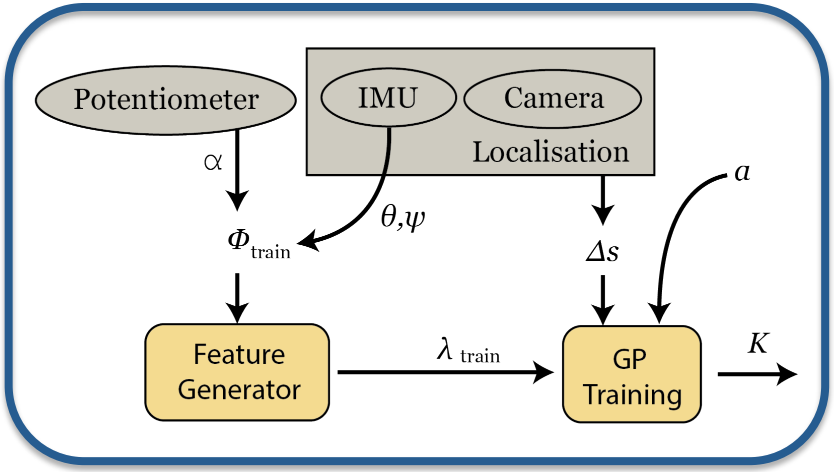

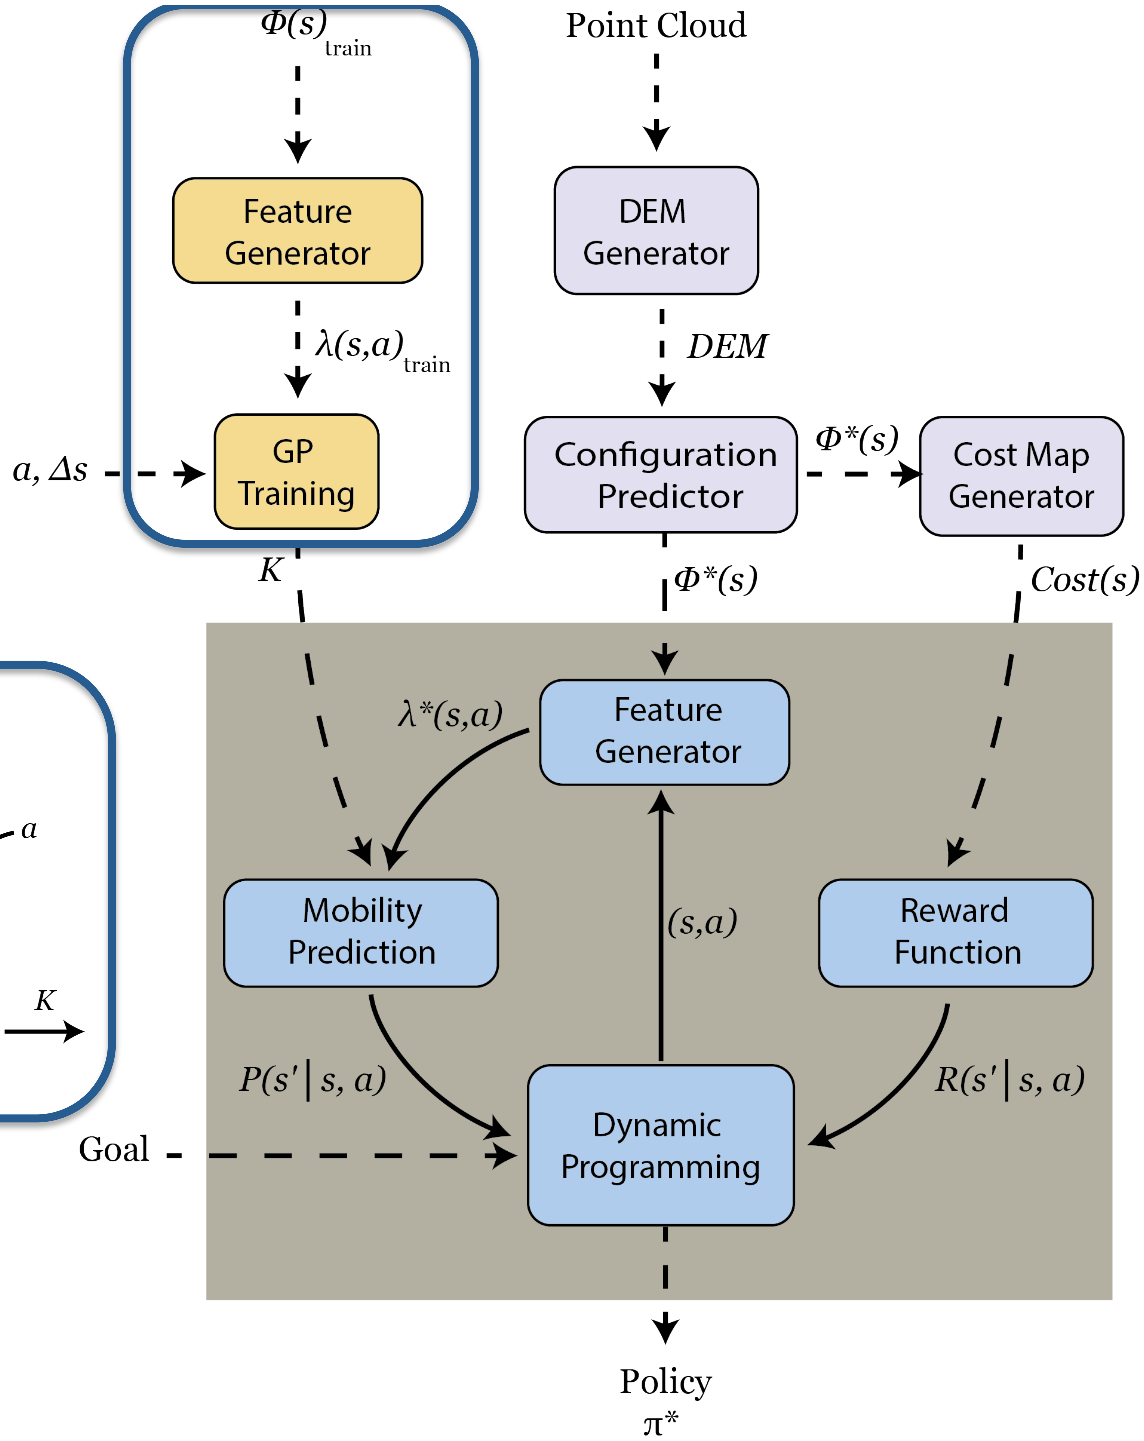

Training

- Multiple executions of each action (total 500+) on large variety of terrain profiles
  - Logging localisation, vehicle attitude & configuration, DEM
- Features λ better representing the terrain profiles (evolution of vehicle attitude and chassis configuration during action execution), determined using PCA on the training data

$$p(\Delta s | \boldsymbol{\lambda}(s, a), a)$$

LfP - Training Data

Training data: heading outcomes (radians) marginalised by action

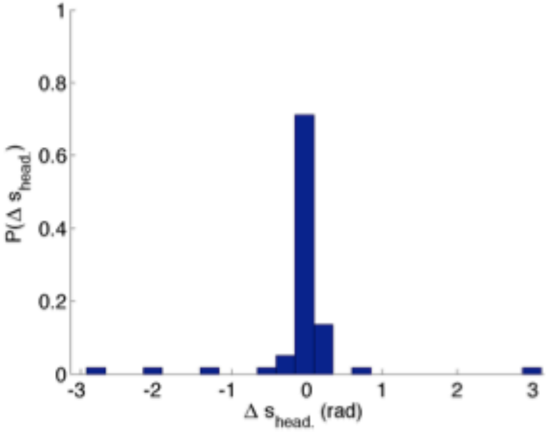

(a) Δs_{head} for crab(0π)

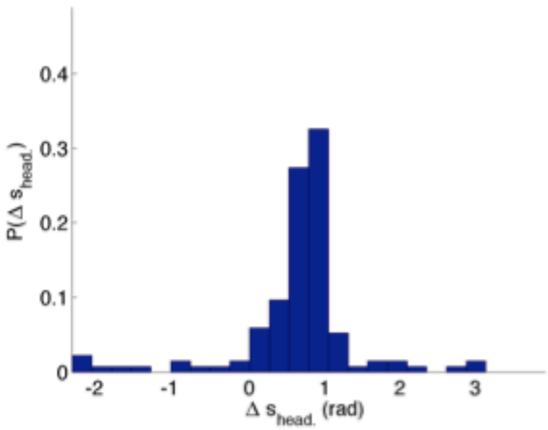

(b) Δs_{head} for crab($\pm\frac{\pi}{4}$)

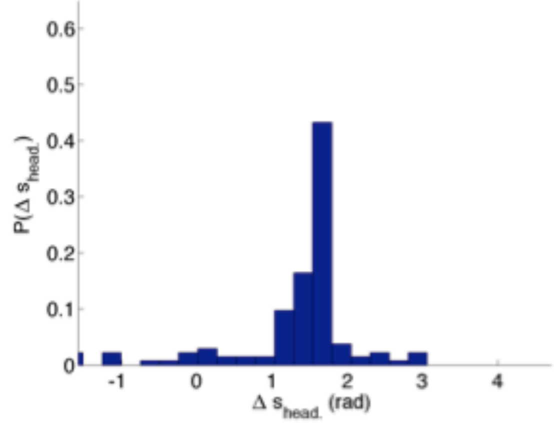

(c) Δs_{head} for crab($\pm\frac{\pi}{2}$)

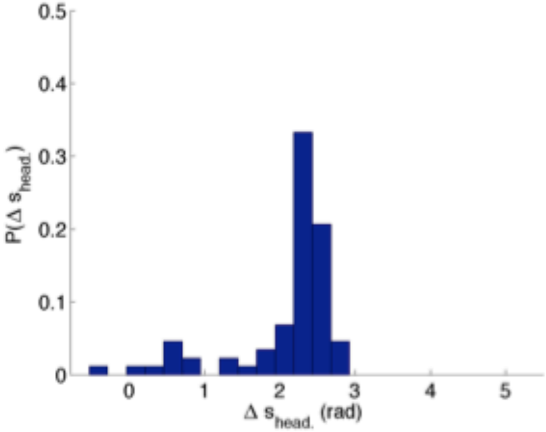

(d) Δs_{head} for crab($\pm\frac{3\pi}{4}$)

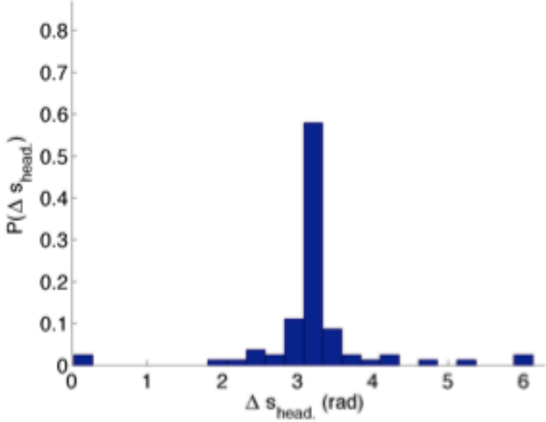

(e) Δs_{head} for crab(π)

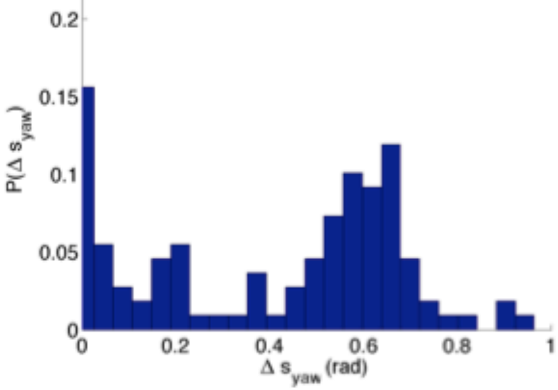

(f) Δs_{yaw} for rotate($\pm\frac{\pi}{4}$)

LfP - Training Data (2)

Training data: distance travelled (m) marginalised by action

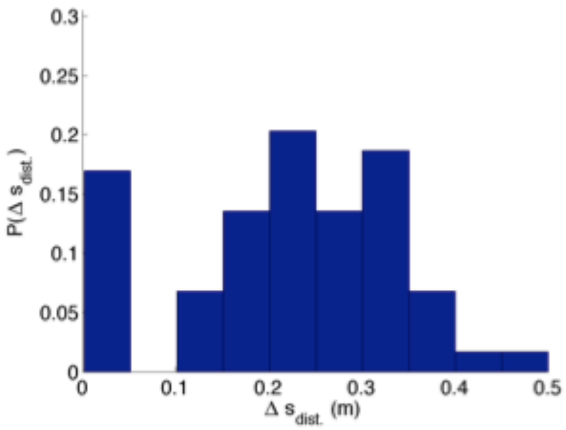

(a) Δs_{dist} for crab(0π)

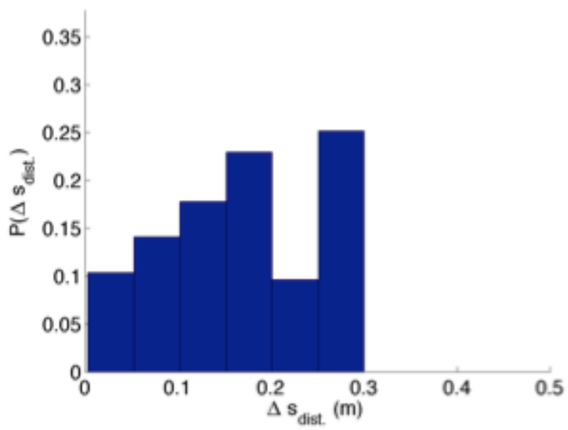

(b) Δs_{dist} for crab($\pm\frac{\pi}{4}$)

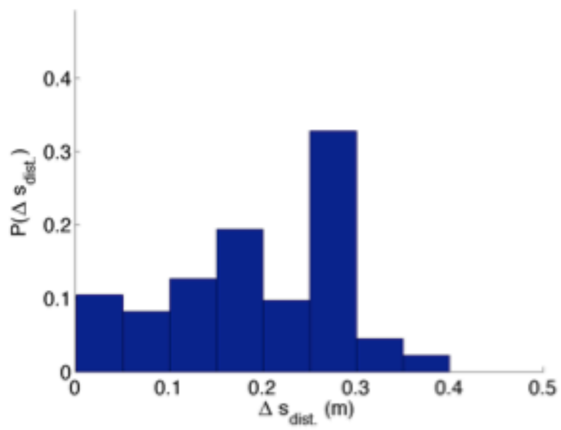

(c) Δs_{dist} for crab($\pm\frac{\pi}{2}$)

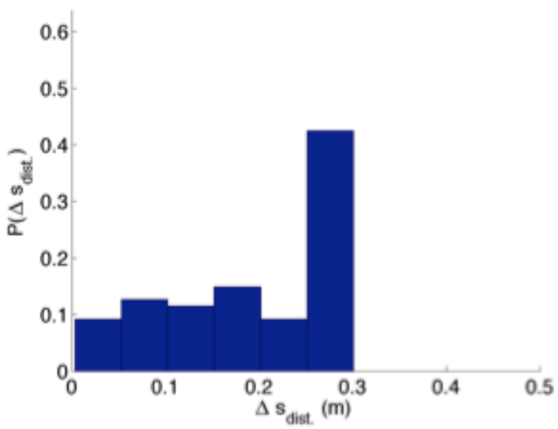

(d) Δs_{dist} for crab($\pm\frac{3\pi}{4}$)

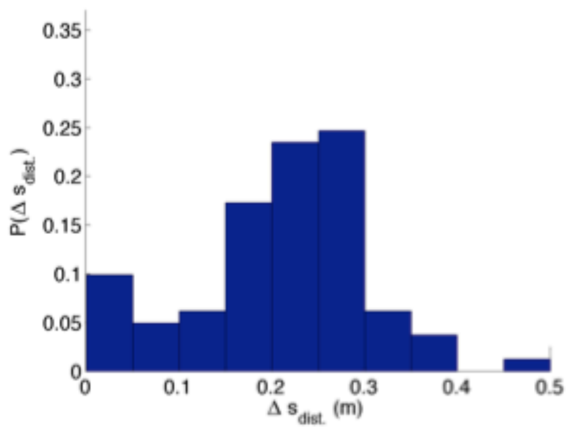

(e) Δs_{dist} for crab(π)

Example of Policies

Policy yaw = 0

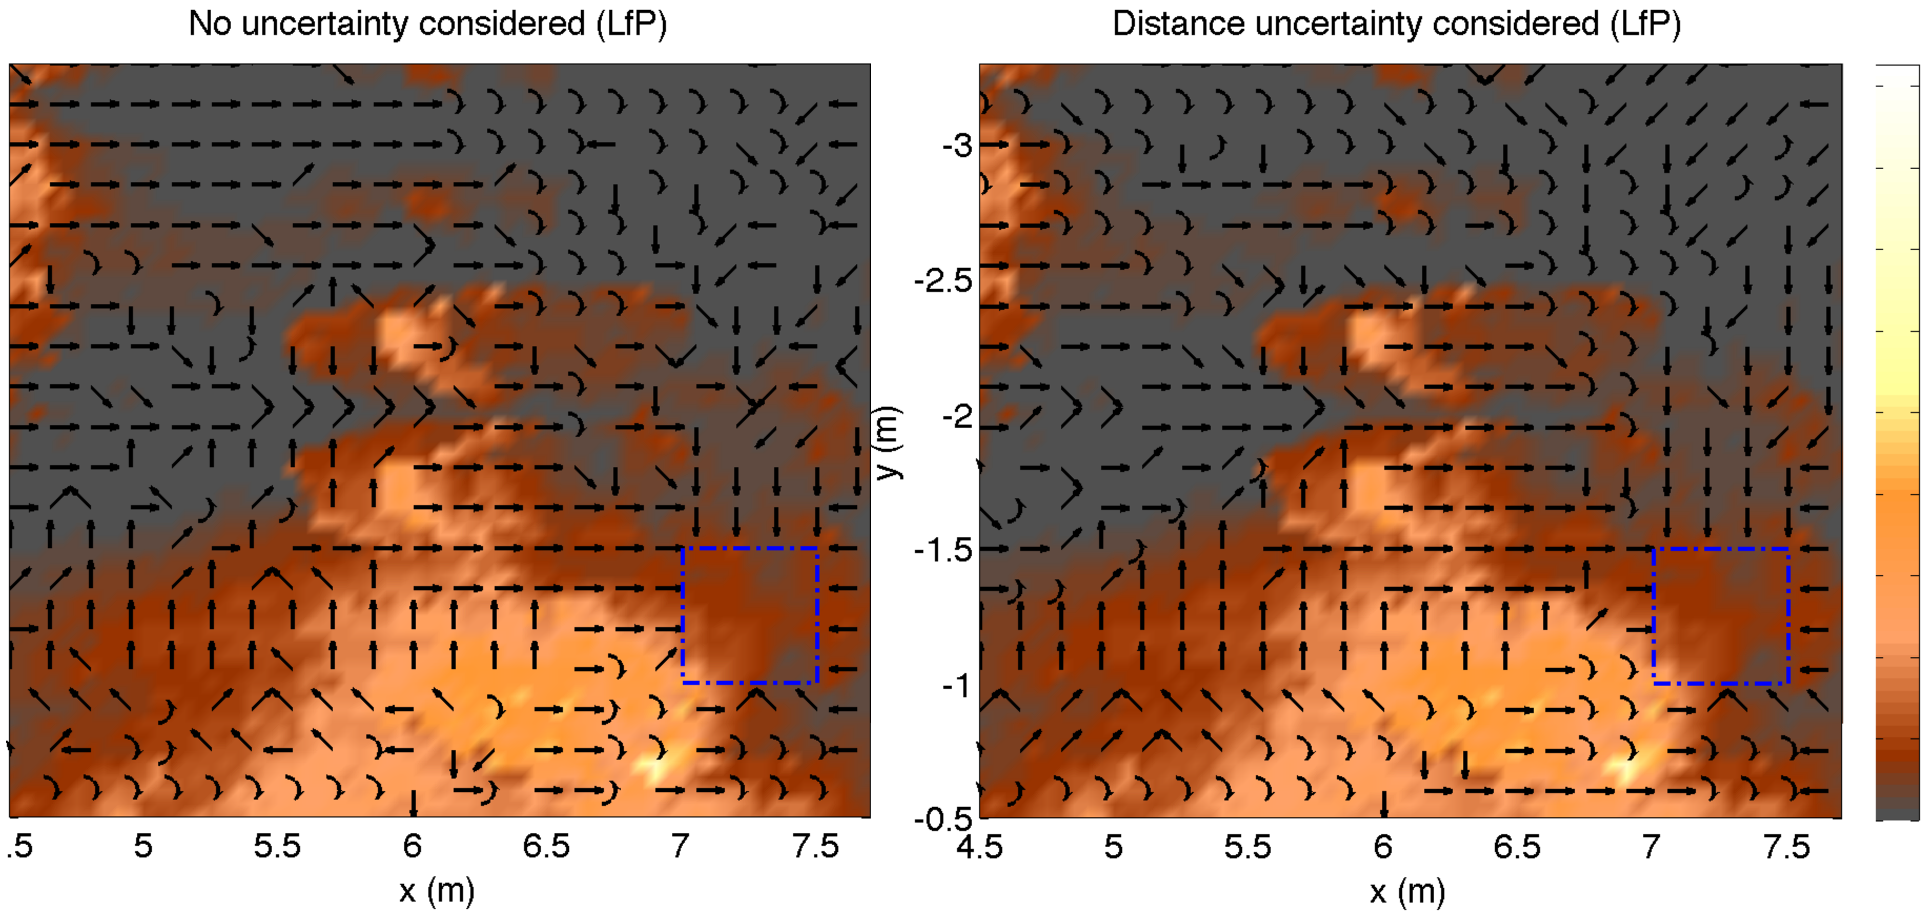

Results: Rigid Terrain Traversal

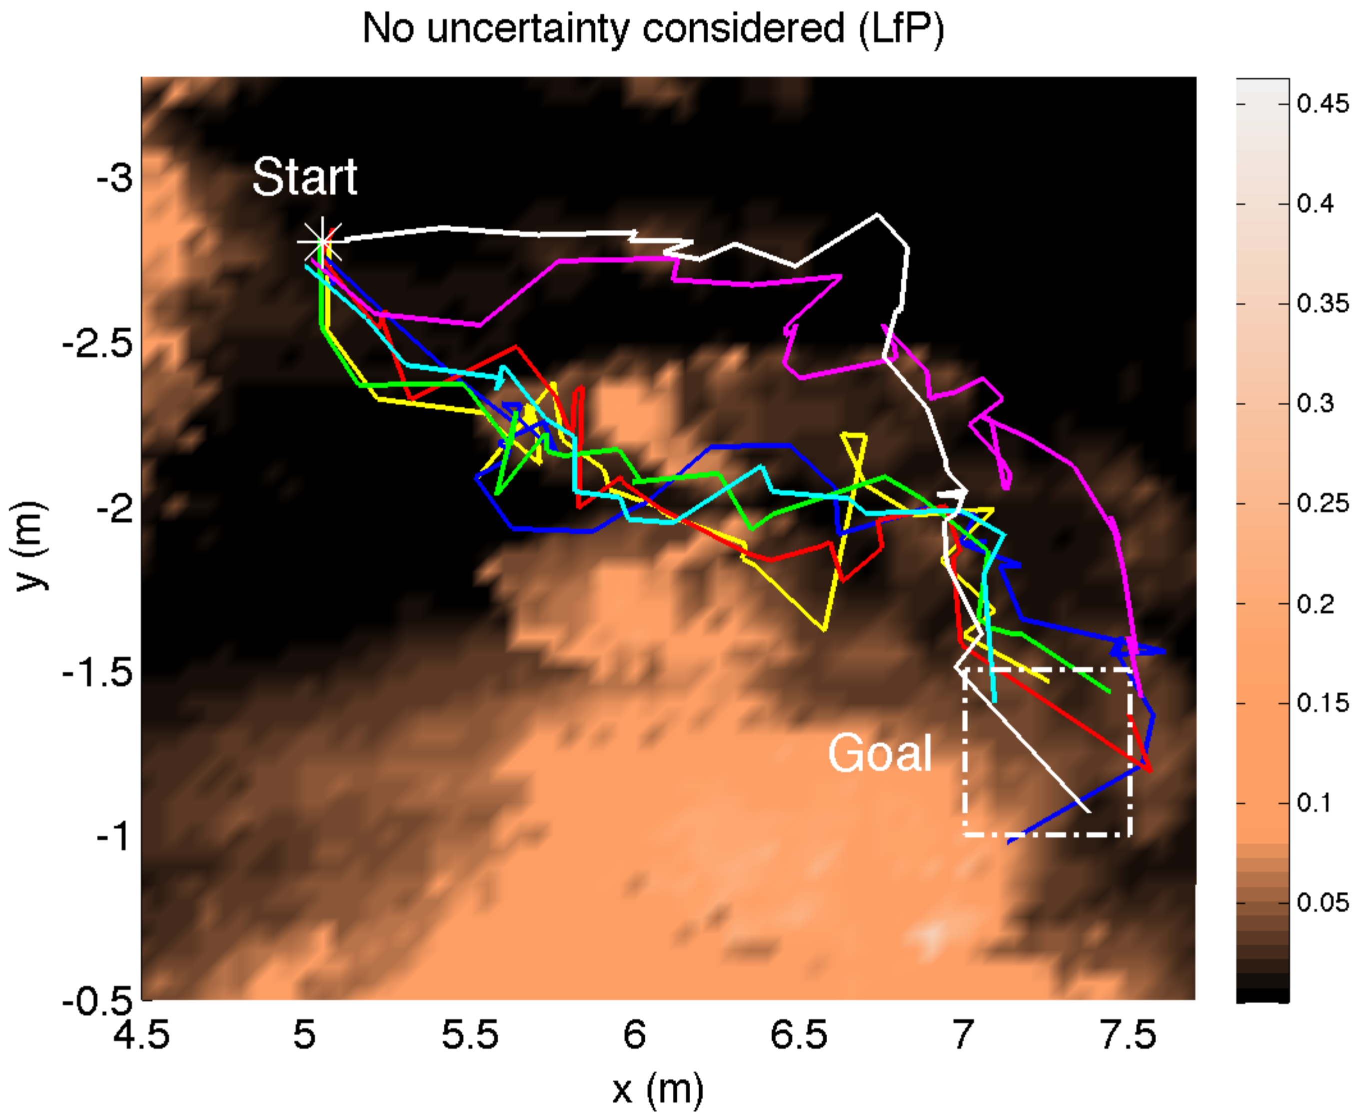

Results: Rigid Terrain Traversal

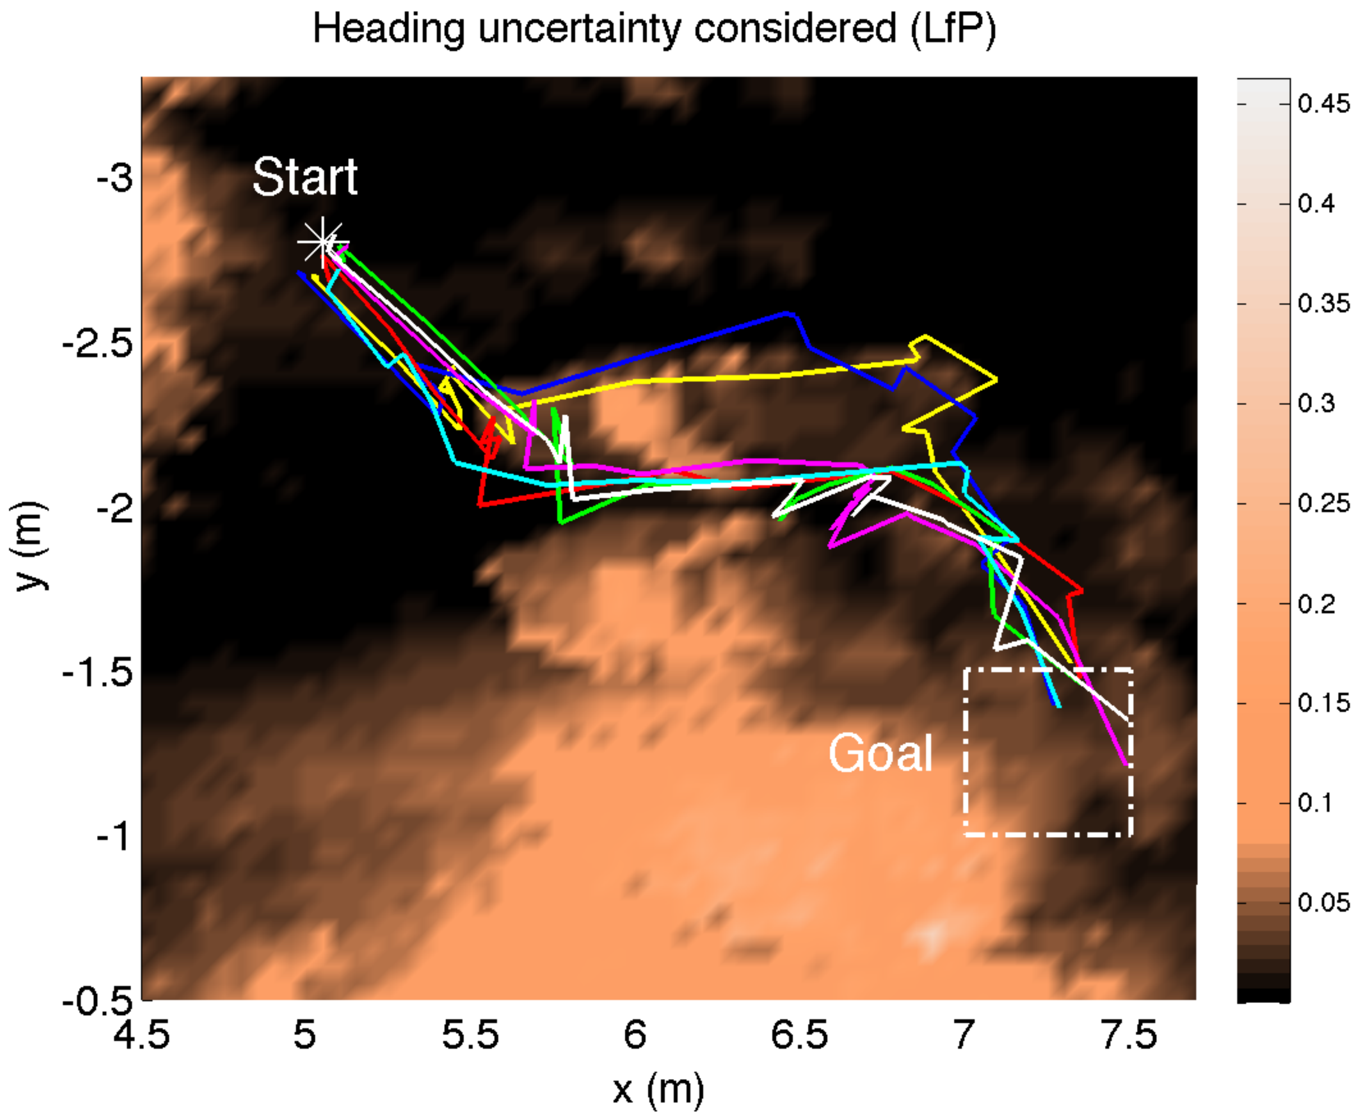

Results: Rigid Terrain Traversal

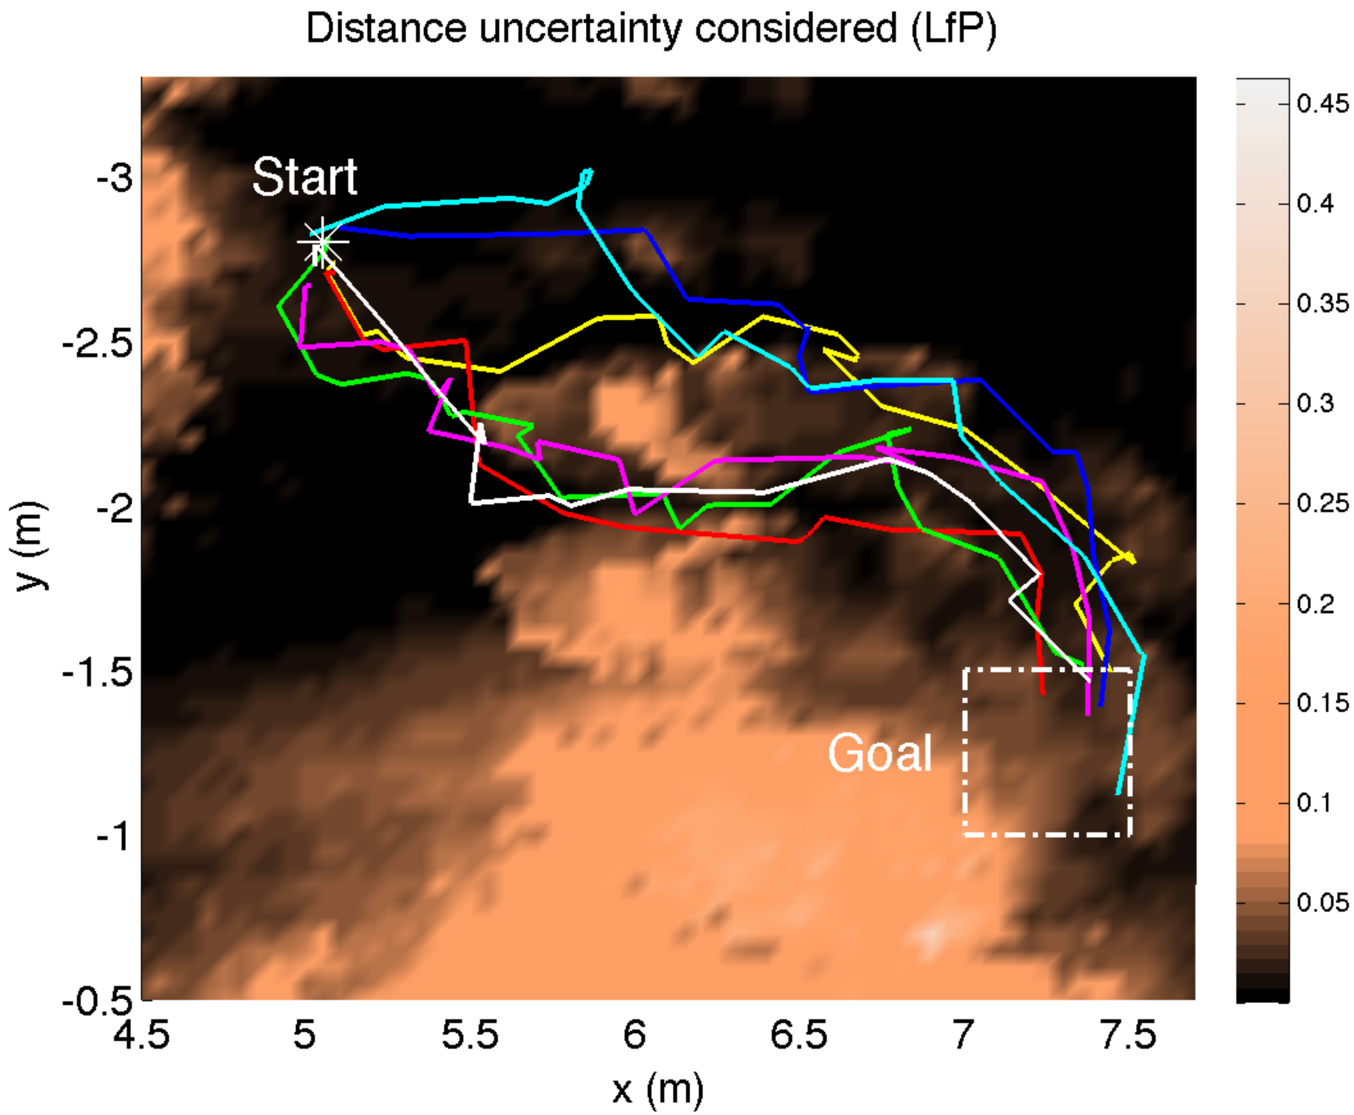

Results: LfP - Rigid Terrain

Uncertainty considered	Total runs	Successful runs (temporarily stuck)	Stuck (%)
None	24	17 (12)	7 (29%)
Distance & Yaw	24	17 (8)	7 (29%)
Heading & Yaw	21	17 (8)	4 (19%)

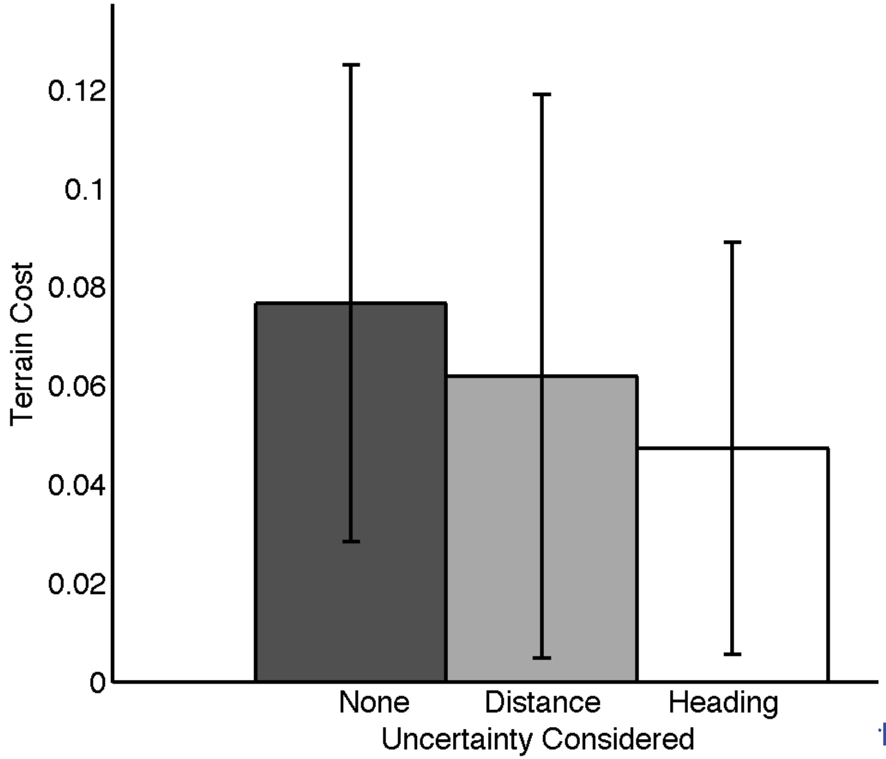

'ksh

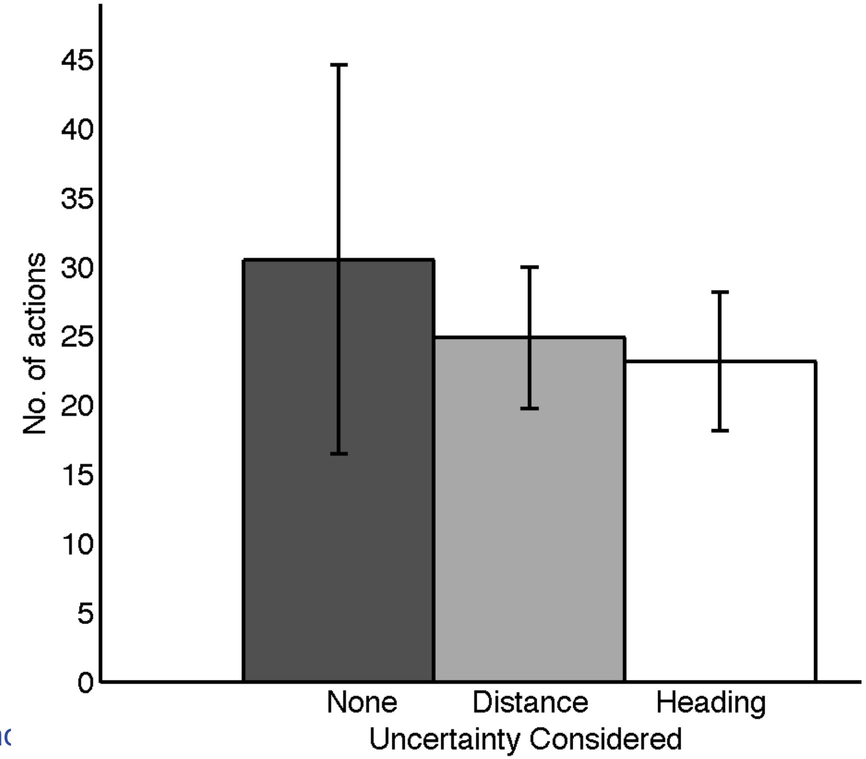

Results: LfP - Deformable Terrain

Uncertainty considered	Total runs	Successful runs (temporarily stuck)	Stuck (%)
None	19	17 (7)	2 (10%)
Distance & Yaw	20	17 (5)	3 (15%)
Heading & Yaw	19	17 (6)	2 (10%)

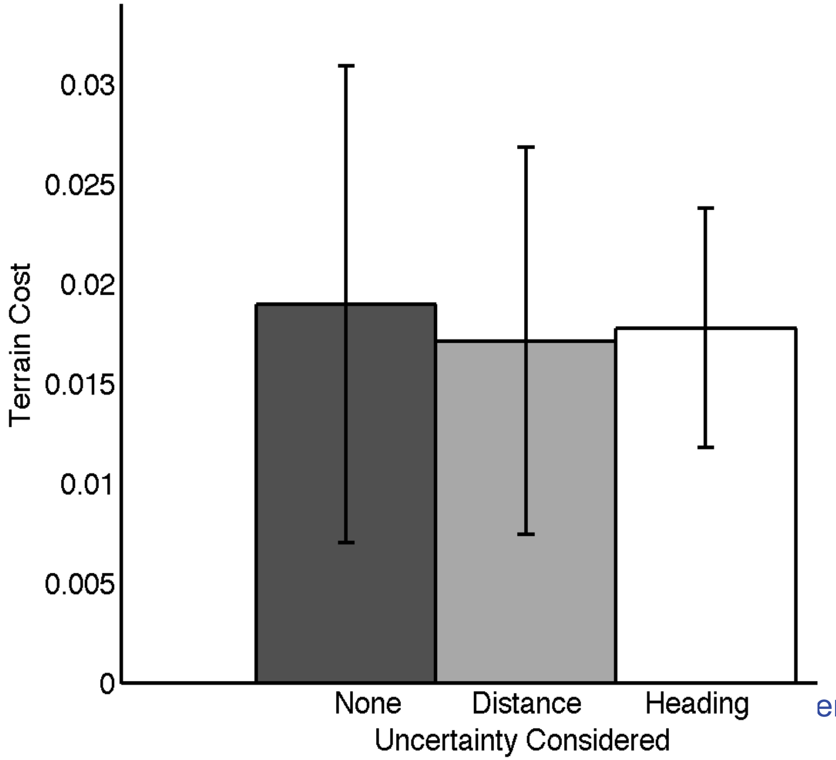

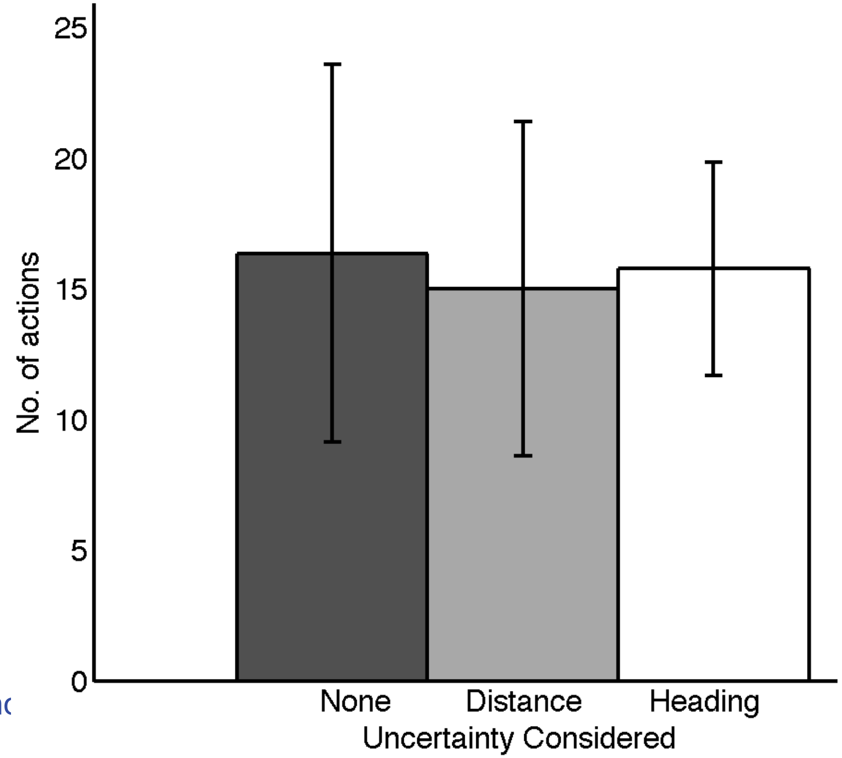

LfP: Conditions/Limitations

- DEM (sufficiently) accurate
  - (accuracy vs. resolution of DEM/state space)
- Correct kinematic model
  - Low speeds
- Terrain *before* rover traversal (as seen by exteroceptive sensors, i.e. DEM) corresponds to terrain *during* rover traversal (when observed via proprioception)
  - i.e. no terrain defomation (rigid terrain)

ACFR
AUSTRALIAN CENTRE FOR FIELD ROBOTICS

Learning Mobility Prediction from Exteroception (LfE)

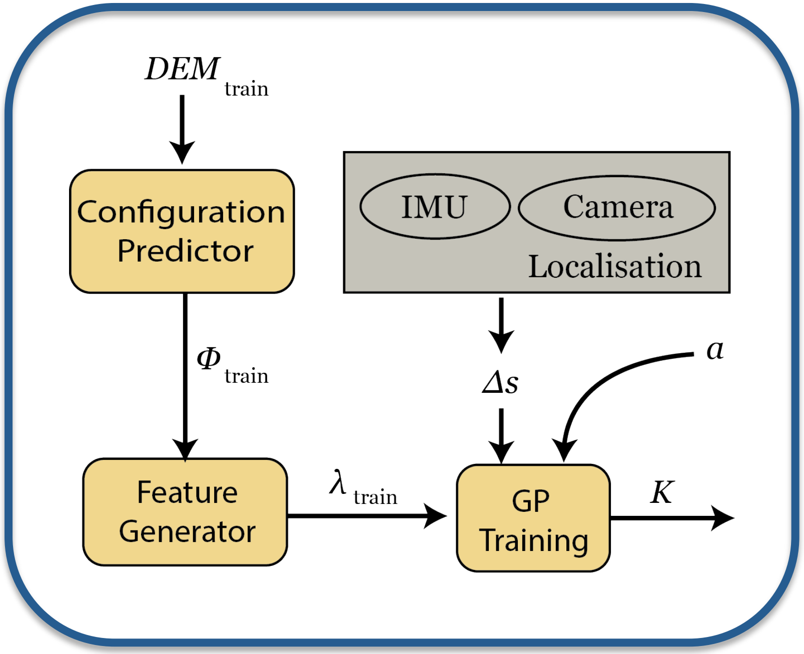

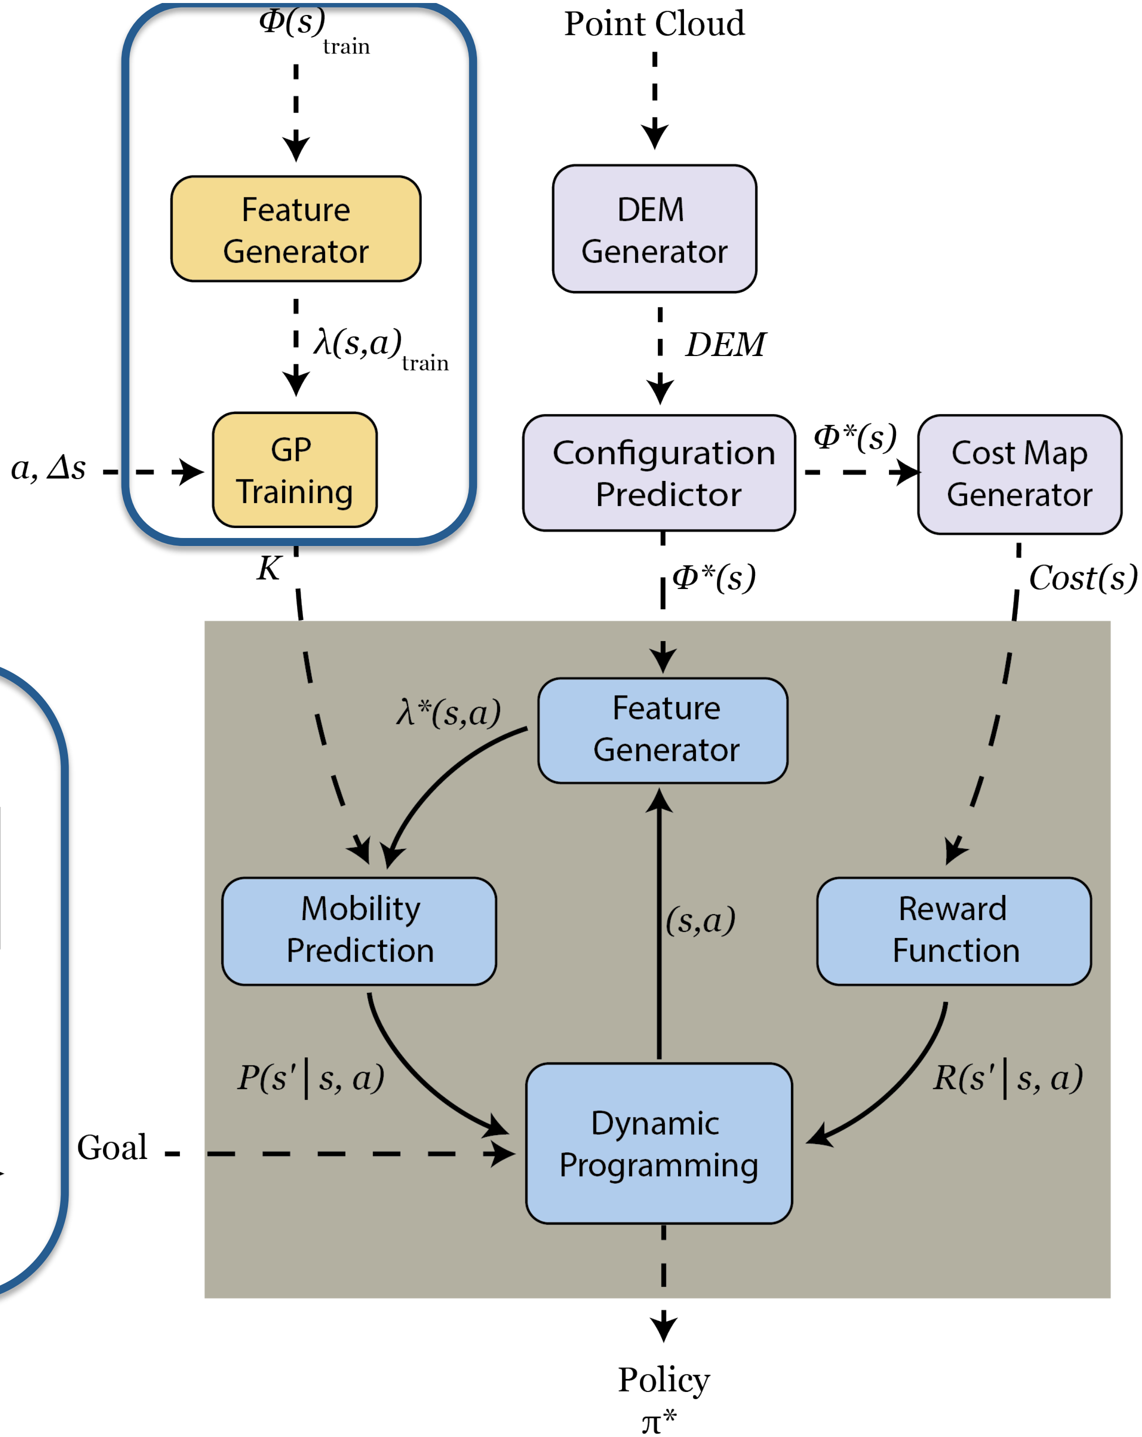

LfP (reminder)

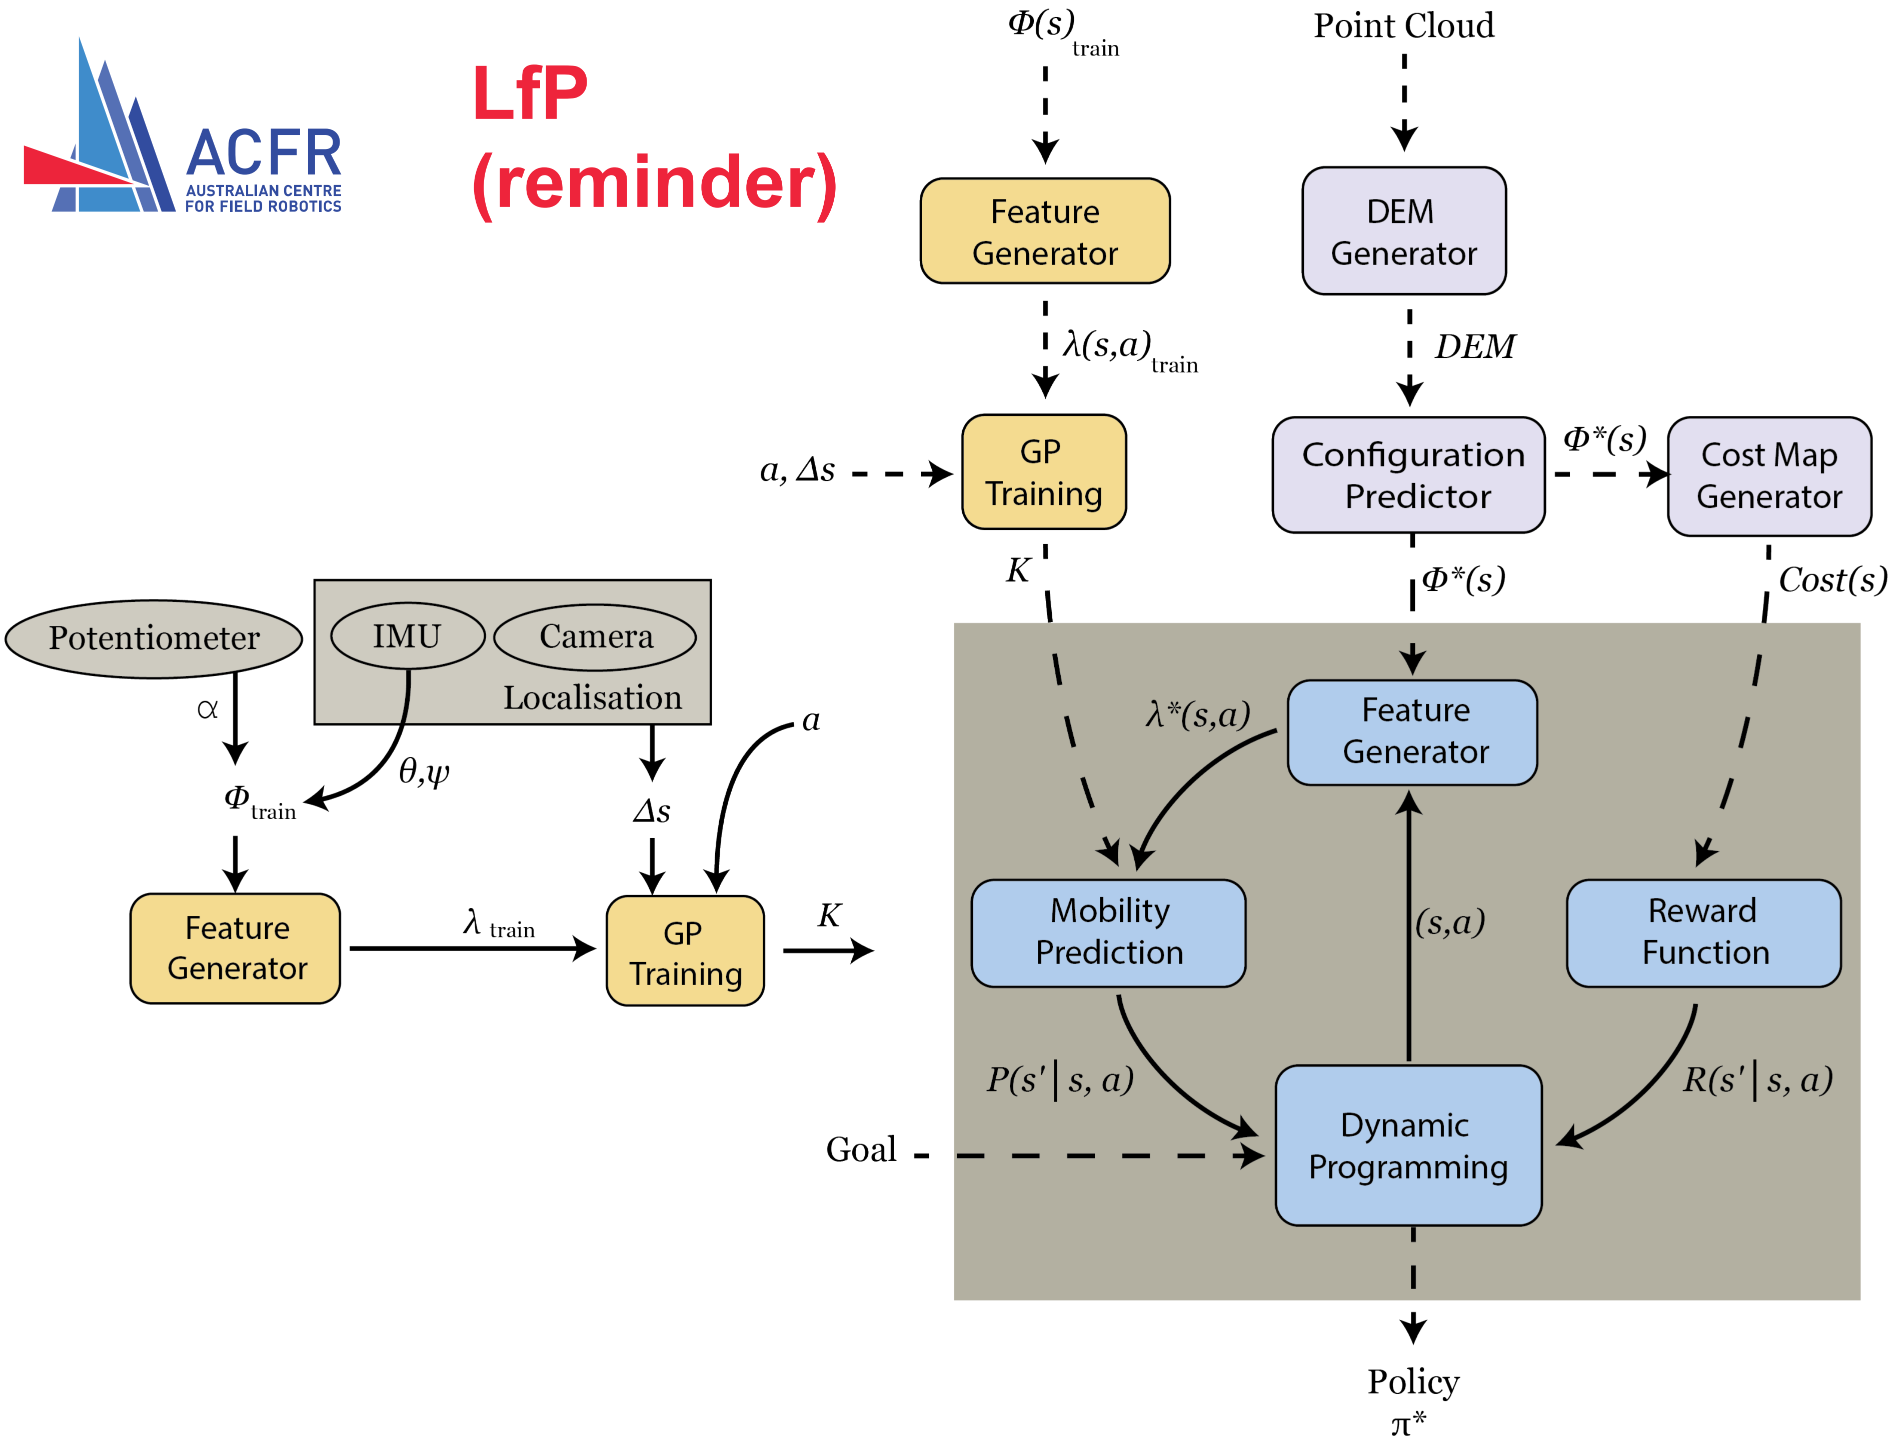

Results: LfE – Rigid Terrain

Uncertainty considered	Total runs	Successful runs (temporarily stuck)	Stuck (%)
None	22	17 (8)	5 (23%)
Distance & Yaw	19	17 (5)	2 (10%)
Heading & Yaw	19	17 (2)	2 (10%)

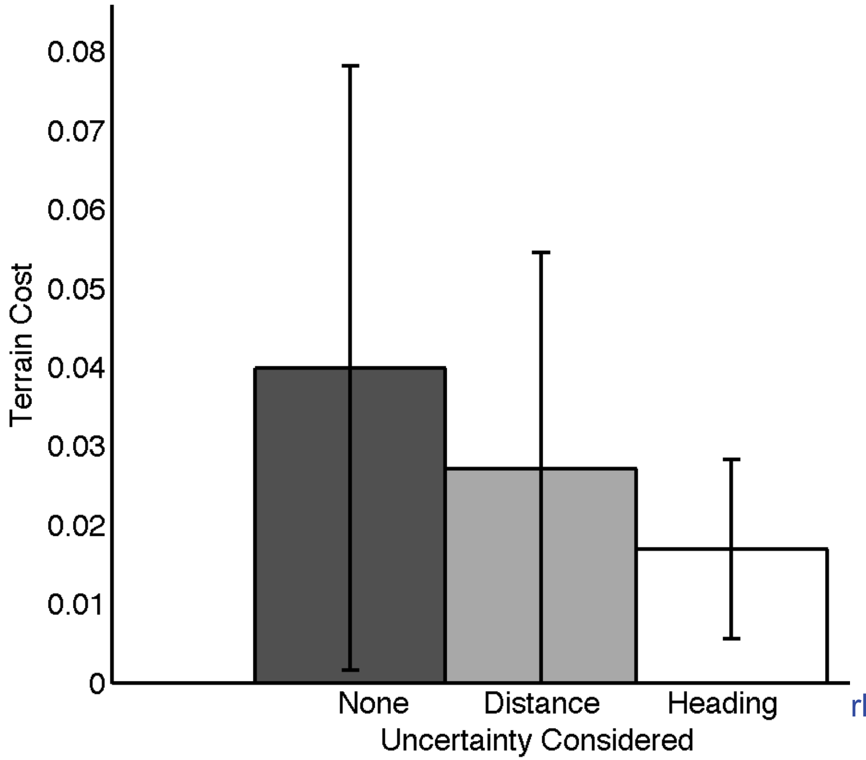

rkshop

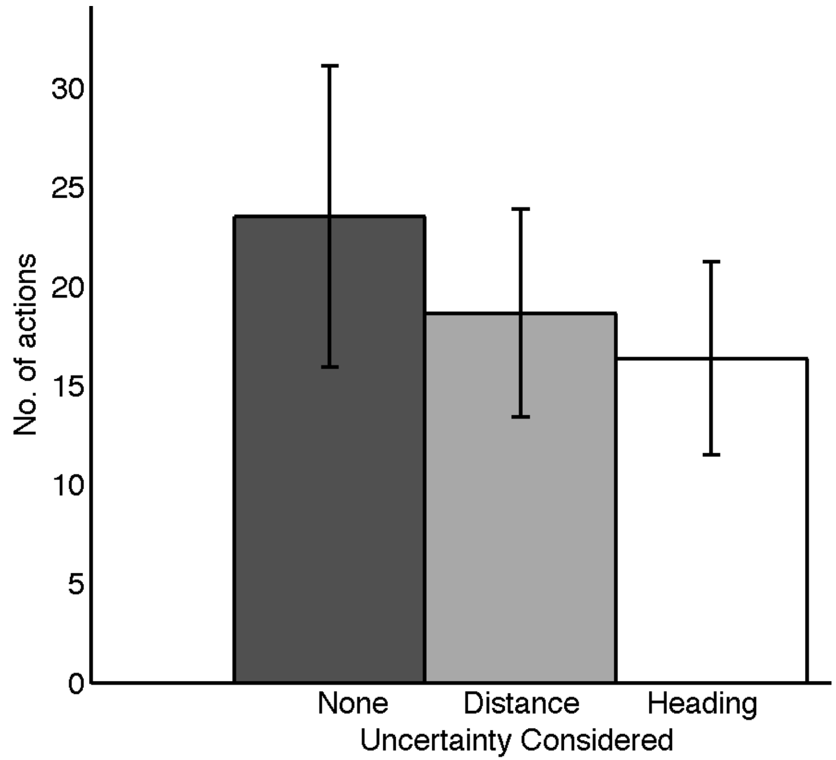

Results: LfE – Deformable Terrain

Uncertainty considered	Total runs	Successful runs (temporarily stuck)	Stuck (%)
None	20	17 (6)	3 (15%)
Distance & Yaw	22	17 (5)	0 (0%)
Heading & Yaw	21	17 (4)	0 (0%)

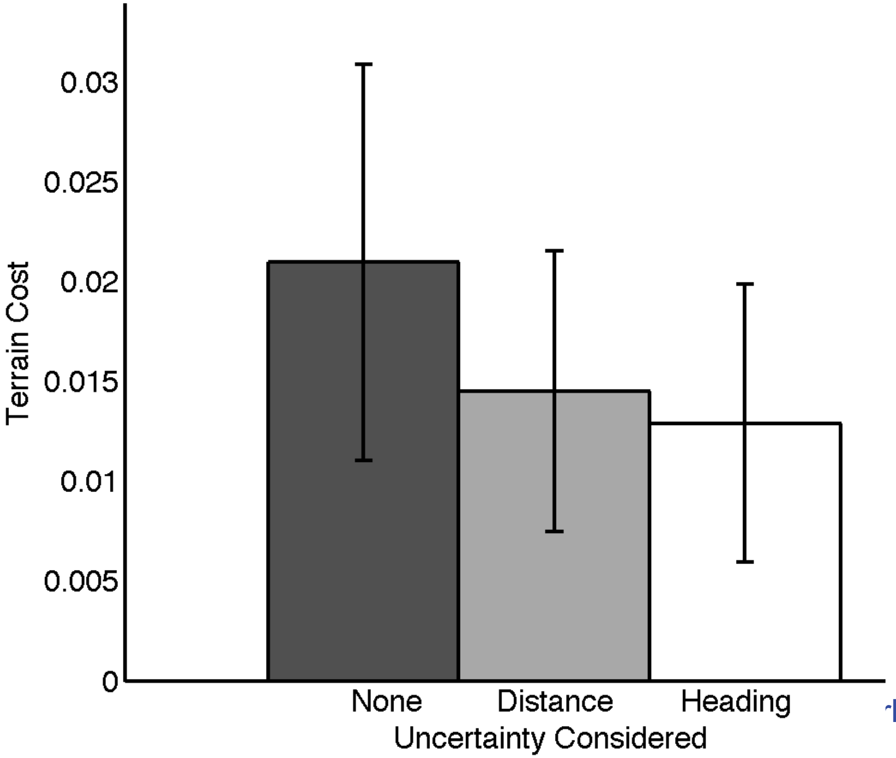

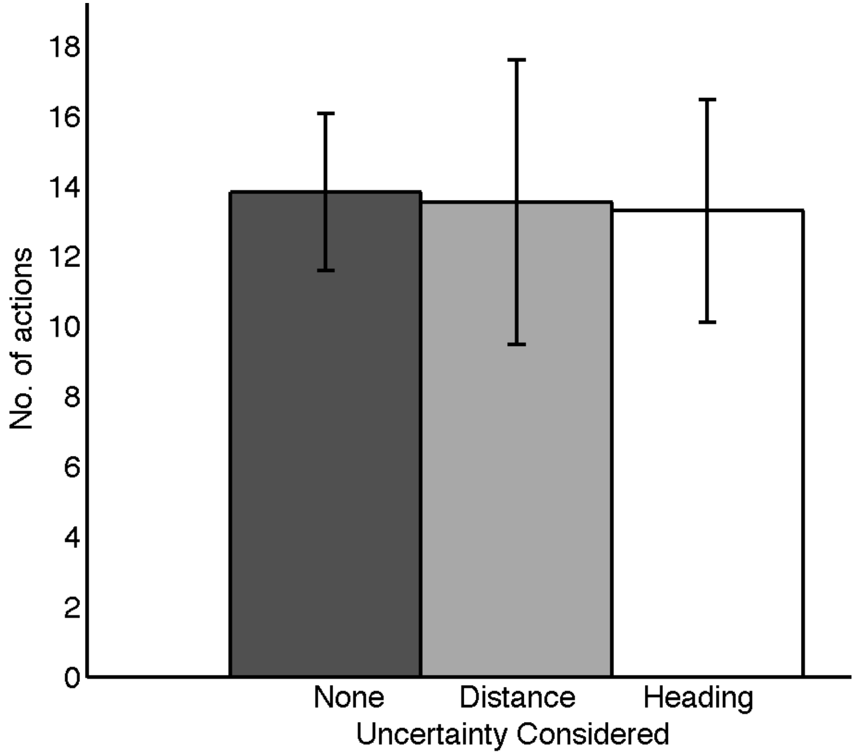

Motion Planning with Stochastic Control

Summary:

- The proposed approach learns a model of control uncertainty directly from experience, which is used explicitly in the computation of a motion policy.
- Experimental validation shows increased reliability (reduced failures) and safety (reduced cost)

Future Work:

- Explicitly account for terrain deformation
- Integrate other types of uncertainty (perception, localisation...)
- Online GP learning & real-time DP (policy updates)

Acknowledgments

The work was supported by the:

- Australian Research Council (ARC)
- Australian Space Research Program Pathways to Space: *Empowering the Internet Generation*
- Air Force Research Laboratory under agreement FA2386-10-1-4153
- Australian Centre for Field Robotics (ACFR)