Finding the ground hidden in the grass: traversability estimation in vegetation

Patrick Ross, Andrew English, David Ball, Ben Upcroft, Peter Corke
ARC Centre of Excellence for Robotic Vision
School of Electrical Engineering and Computer Science
Queensland University of Technology
patrick.ross@connect.qut.edu.au

Abstract

This paper presents a novel online algorithm for estimating ground height in vegetated environments, and its application to traversability estimation. Traversability estimation based on the shape of the ground is difficult in many environments where the ground is partially to completely obscured by vegetation. Typical algorithms assume that the ground is directly observable with no noise, which leads to incorrectly estimating the ground height and hence traversability in the presence of vegetation. The algorithm learns the typical distribution of visually observed heights around the ground, and probabilistically infers the true ground height. It iteratively improves its internal estimate over multiple frames. The iterative nature allows the algorithm to run in limited processing time, and to become more accurate with increased processing time. The results demonstrate that this algorithm out-performs other strategies in estimating ground height, and can be used to reduce the vertical acceleration of the vehicle in real time.

1 Introduction

Robotic platforms operating in outdoor environments require the ability to robustly perceive potential dangers beyond just binary obstacle detection. This includes the ability to observe and adjust the robot's motion for the underlying terrain through traversability estimation. The type, shape and continuity of the terrain contribute to its traversability, which influences the actions the robot must take to traverse the terrain safely; a robot should slow down for irregular or bumpy terrain. Each vehicle will have a particular tolerance for maximum and accumulated stresses due to peak accelerations.

Most techniques assume that the ground is fully observable, and operate under the assumption that all observed points belong to the ground [Berczi *et al.*, 2015; Goldberg *et al.*, 2002; Ho *et al.*, 2013]. This assumption is true in environments such as deserts and man-made environments. In unstructured environments, such as many field environments, features such as grass and other plants will partially or completely obscure the ground, and therefore simple approaches will have significantly degraded performance. Additionally, it is often assumed that the traversability of a patch of ground is independent of the direction of traversal, which is an invalid assumption for many of these common environments. These factors lead to an incorrect estimation of traversability, limiting the utility of the robotic platform. This work aims to expand the domain of environments where traversability can be estimated to include grassed environments and those with low-lying vegetation, such as agricultural fields.

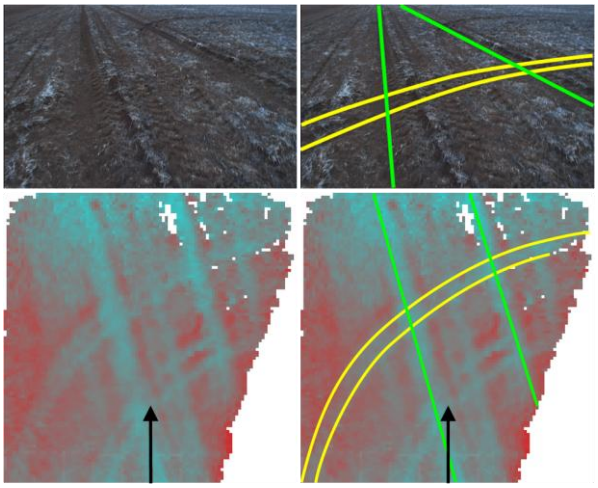

Figure 1: Typical headland in a broad-acre field. The wheat stubble present obscures the ground, corrupting any measurements of the true ground location. This results in the crop dominating traversability measures, as opposed to the true danger, which is the deep wheel tracks. The first row shows camera imagery, the second indicates the estimated ground height (blue is low, red is high) from the proposed algorithm, with the robot indicated by the arrow. The right figures indicate correspondences between the two.

This paper presents a novel online algorithm for estimating the location of the ground even in environments where it is partially or completely obscured. The algorithm takes periodic observations of the typical distribution of visually observed heights around the ground, and probabilistically infers the location of the ground given noisy measurements. It makes no assumptions about the distribution of heights, and so is equally applicable to many environments, including those with no vegetation. It learns online and therefore is capable of adapting to changing environmental conditions. The results demonstrate that this algorithm significantly out-performs other strategies for ground height estimation in complex

field environments.

This algorithm is demonstrated operating online, regulating the speed of a mobile robot platform to reduce its peak accelerations. The ground height estimation algorithm is combined with a path-dependent traversability metric that computes the maximum applicable velocity with respect to the shape of the ground. The results demonstrate that the robot slows down for hazards in the environment, reducing peak accelerations on the vehicle.

The remainder of this paper is laid out as follows. Section 2 provides relevant background to the work presented in this paper. Section 3 outlines the approach. Sections 4 and 5 introduce the experiments and detail results respectively, while Section 6 provides conclusions.

2 Background

There are two different approaches to traversability estimation: 1) those that estimate traversability directly from sensor measurements, typically using some form of machine learning, and 2) those that model important characteristics of the environment and predict the attitude of the vehicle as it traverses these regions.

Approaches that learn traversability directly from sensor measurements require either provided learning inputs from a human supervisor, or interactions with the environment for example using a bump sensor [Kim *et al.*, 2007; Kim *et al.*, 2006] or by wheel slip estimation [Angelova *et al.*, 2007]. Pre-trained strategies are limited in their generalization to different environments. Learning via interaction means that the robot must interact with potentially dangerous environmental features before it learns to avoid them, and as such its safety cannot be guaranteed.

Strategies which are based on environmental and robot characteristics are capable of generalizing well to different environments, provided the important characteristics can be well determined. Goldberg [Goldberg *et al.*, 2002] introduced the GESTALT traversability metric, which uses plane fit statistics to determine traversability, together with known information about the robot such as its minimum clearance. Other works determine traversability by estimating the attitude of the robot on a given surface [Ho *et al.*, 2013; Lacroix *et al.*, 2002; Tarokh and McDermott, 2005], and inferring statistics such as the stability margin. These strategies assume the load bearing surface can be directly observed by taking the median of observations in a vertical cell, and as such fail in environments where the ground is obscured.

Typical traversability algorithms operate in close field of view where good structural information is available. Traversability estimation can be extended to a longer range using near-to-far learning [Hadsell *et al.*, 2009; Sofman *et al.*, 2006; Vernaza *et al.*, 2008].

Wellington [Wellington *et al.*, 2005] proposed a method to infer the location of the ground in environments where the ground is partially or completely obscured. This method classified observations into ground or vegetation classes, and then attempted to ascertain a common vegetation height to predict the likely location of the ground. This method assumes that all vegetation on the ground is of a similar height, which is an invalid assumption in many environments, and additionally requires highly dense range measurements and multispectral imagery to accurately classify and infer from observations.

3 Design

The basis of the algorithm is the fact that observations of heights are drawn from some underlying distribution, which can be approximated. The method is iterative, and so with more observations the results become increasingly accurate.

The algorithm has three main stages as shown in Figure 2. Firstly, the distribution of heights is estimated relative to the location of the ground, denoted model construction. The model construction module maintains a combined point cloud in the current reference frame, from which it constructs the model. Secondly, this model is used to iteratively infer the location of the ground in noisy and incomplete data. The ground height inference module iteratively resamples the ground height estimate using the constructed model as a reference point, and combines multiple samples over time to produce an estimate of the true ground height.

Finally, this model of the ground is used to estimate the traversability of the terrain given planned path. This is achieved indirectly by determining the maximum applicable velocity of the robot given a required maximum vertical acceleration.

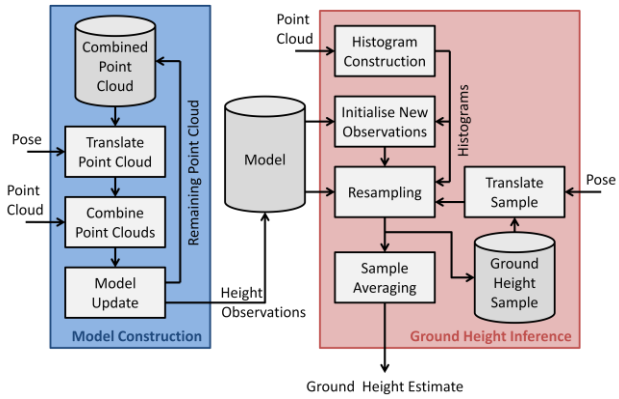

Figure 2: Ground height estimation pipeline, indicating the two major components - model construction and ground height inference.

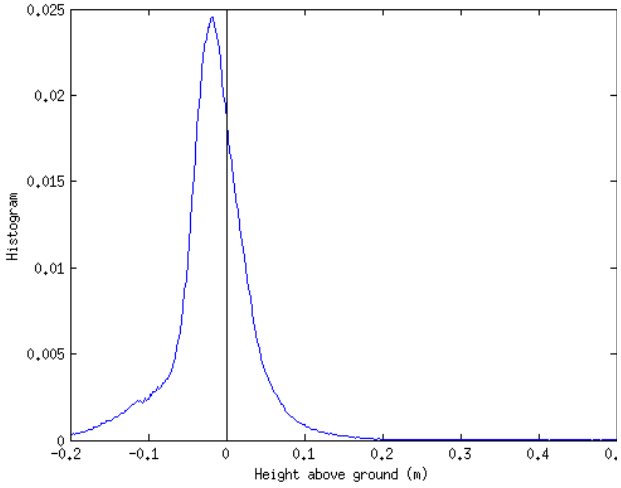

Figure 3: Typical model histogram for stubble fields showing the distribution of height observations around the true ground. The histogram includes noisy stereo matching results and stubble.

3.1 Model construction

The model is a histogram which describes the probability of observing a particular height relative to the mechanically observed ground height. The histogram is built online which enables the system to adapt to environments with different types of vegetation.

The histogram is constructed by tracking points until they pass underneath one of the robot's wheels, where an estimate of the ground height can be inferred from mechanical constraints. This gives the absolute location of the points relative to the ground, which can then be added to the model. Samples are discarded when they move behind the robot. Old samples are slowly forgotten to allow for the model to handle transitions between environments.

Figure 3 shows a sample histogram for the environment investigated in this work. In this example, the nonzero regions below -0.1m are due to errors in the stereo matching algorithm. The peak is slightly misaligned with the true zero. This is because the true location of the bottom of the wheel, the location used for absolute grounding of the measurements, is a compliant surface. This misalignment has no effect on the results in the camera's reference frame, since it is effectively accounting for the errors in other transformations. The slow drop off in probability above the ground is a result of the stubble (crop residue) present in the environment.

A consequence of this model construction strategy is that the system can only learn about terrain that the robot drives over. In some circumstances the robot could avoid all vegetation, making the model heavily biased.

3.2 Ground height inference

The ground height is inferred by assuming that all height samples were drawn from the model. Provided the robot has traversed a sufficiently representative region of the world, this is a valid assumption.

To infer the ground height, the world is first discretised into a grid of 100×100 cells, each of which is 0.1m × 0.1m. The size of the region was limited for two reasons; firstly computation savings, and secondly outside this region the accuracy of stereo matching was significantly decreased. During each update step, a histogram of heights is determined for each cell from the current point cloud. The algorithm iteratively resamples the ground height of each cell using the conditional probability

$$P\left(h_i \middle| h_{N_i}, O_i\right) \propto P\left(h_{N_i} \middle| h_i\right) P(O_i|h_i)\, P(h_i) \qquad (1)$$

where h_i is the height of cell i, h_{N_i} is the average height of the 8-connected neighbors of cell i, and O_i is the histogram associated with that cell. Resampling each height in turn from the conditional probability tends towards the true distribution according to Gibbs sampling [Geman and Geman, 1984]. The benefit of this strategy is that resampling can be time limited per update, and continue to improve with further updates.

Periodic samples are averaged to provide an estimate of the true ground height for each of the cells.

This conditional probability consists of three components; the smoothing component, the data component and the prior. The prior is defined as

$$\log P(h_i) = -\lambda_g h_i \qquad (2)$$

where λ_g is a system parameter to bias the solution towards lower ground heights, which is necessary for situations where the ground may be completely obscured. The data component is defined as

$$\log P(O_i|h_i) = -KL\left(O_i \middle| M(h_i)\right) \qquad (3)$$

where $M(h_i)$ indicates the model histogram shifted by the given height, and $KL(X|Y)$ is the Kullback-Leibler divergence between the two histograms. This probabilistically makes solutions tend towards aligning the observations with the model. The smoothing component is defined as

$$\log P\left(h_{N_i} \middle| h_i\right) = -\frac{\left(h_i - h_{N_i}\right)^2}{2\sigma_g^2} \qquad (4)$$

where σ_g is the standard deviation of the ground height, and is also a system parameter.

The condition probability $P\left(h_i \middle| h_{N_i}, O_i\right)$ is sampled using the Metropolis-Hastings algorithm [Hastings, 1970; Metropolis *et al.*, 1953], using a Gaussian proposal distribution. Resampling using the Metropolis-Hastings algorithm requires an initial guess to be close to the true answer, otherwise the resampling can take an arbitrary amount of time to converge to the true distribution. To this end new samples are approximately initialized using a brute force strategy.

3.3 Traversability estimation

The traversability of ground regions is dependent on the path taken by the robot through the environment; hence the traversability of the ground can only be computed once the path is known.

The path of each of the wheels through the ground height map is calculated using their known offset from the robot. The ground height at each of the points along this path is calculated.

To calculate the maximum velocity, it can be shown that external accelerations (due to sources other than the vehicle's propulsion) are given by

$$a = \left[1 + \left(\frac{dy}{dx}\right)^2\right]^{-\frac{3}{2}} \frac{d^2y}{dx^2} v^2 \qquad (5)$$

where x, y describe the shape of the ground, v is the velocity of the robot tangential to the ground, and a is the acceleration magnitude. It is assumed that the velocity perpendicular to the ground is zero.

Based on this, a maximum velocity can be computed, given as

$$v_{max} = \left[1 + \left(\frac{dy}{dx}\right)^2\right]^{\frac{3}{4}} \sqrt{\frac{a_{max}}{\frac{d^2y}{dx^2}}} \qquad (6)$$

At each point along the paths, the maximum velocity for each wheel is determined, and then these are combined by taking the smallest maximum velocity at each point. Acceleration curves are applied to this result to ensure that the minimum velocity requirements can be realistically achieved, and then the new maximum velocity at the current location is applied.

4 Experiments

The platform used for testing is a retrofitted John Deere TE Gator. The vehicle is approximately 2.6m long and 1.5m wide. The cameras used are a stereo pair of Point Grey Grasshopper GS3-U3-23S6-C cameras with a 0.75m baseline, triggered externally at 10Hz. These cameras were mounted at a height of 1.55m, angled downwards 15 degrees. Figure 4 shows the platform.

Stereo calibration was performed using the AMCC

Figure 4: The robotic platform used for testing.

MATLAB toolbox [Warren *et al.*, 2013]. Stereo matching was performed using LIBELAS [Geiger *et al.*, 2011]. Visual odometry was performed using LIBVISO2 [Kitt *et al.*, 2010]. The platform used Robot Operating System (ROS) for path planning and interprocess communication as described in [Ball *et al.*, 2015]. To modulate speed the process intercepted the desired velocity output from the path planner and sent new velocities to the vehicle controller.

Datasets were gathered on a farm in Emerald, Australia. Experiments were carried out in a field with wheat stubble.

4.1 Ground height accuracy

The purpose of this experiment was to evaluate the accuracy of the proposed algorithm against other commonly used methods for ground height, including median, mean and minimum values from the point cloud.

To this end a dataset was gathered of stereo camera imagery in a 10m × 10m region of a field containing predominantly wheat stubble by driving several times through the same region. A series of posts were placed in the field with 5m spacing between to act as robust markers to assist with registering the datasets. The chosen section of field contained deep wheel tracks from farm equipment, as well as some lighter cross wheel tracks.

All of the stubble was then carefully removed from the region so that the ground wasn't obscured at any point. There was little change to the underlying ground height during this process. A second dataset was gathered of this region using the stereo camera pair. This enabled the generation of a high-precision ground truth of the surveyed region, subject to the accuracy of stereo matching and visual odometry. Since the datasets were very short, these were assumed to be locally accurate.

To combine the data into a complete ground truth, the data was split into individual passes of the region which are locally accurate. These were then aligned using a combination of inspection and ICP methods to ensure a high degree of alignment. The ground truth height for each cell was then chosen as the median of the observations as this measurement is the most statistically robust measurement.

For visualization purposes, a second order surface was fitted to the points to remove the slow gradient changes and highlight the local variation. The cells were colored by their deviation from this surface.

4.2 Traversability controller

The purpose of this experiment was to demonstrate the utility of the traversability metric based on the ground height estimator in minimising the acceleration of the vehicle. To this end, a fixed path was mapped out in the field which consisted of sections following the crop rows, travelling at an angle to the crop rows, and travelling directly transverse the crop rows. There was a narrow washout through the area.

Two traverses of this path were performed, one without speed modulation and one with the speed modulation to show the corresponding decrease in the acceleration of the vehicle. Figure 6 shows the path of the robot in each of the traverses, as logged by a high precision GPS-INS system. The raw INS measurements were used to estimate the acceleration of the robot.

5 Results

5.1 Ground height accuracy

Figure 5 shows the predicted ground height for each of the methods. It can be seen that the minimum consistently underestimates the ground height, excepting where crop is present. This is a result of the frame to frame errors which typically are found in a method like stereo matching. Despite this, Wellington [Wellington and Stentz, 2003] demonstrated that even with high accuracy measurements from a laser, the minimum height consistently underestimates the true ground height, making this method unsuitable.

The mean and median give much more robust results, and are comparable, however with a slight reduction in error from the median. It can be seen that the mean is corrupted by erroneous stereo matching results, shown as the deep stripes at a 45 degree angle. The median, on the other hand, is significantly more robust to these outliers.

The median overestimates the height consistently in the middle of the image, corresponding to the significant stubble present in this region.

The proposed method has removed much of the stubble from the output, while maintaining the shape of the ground. Table 1 summarizes the mean error, RMS error and 95th percentile error of the prediction errors for each of the methods. The proposed method can be seen to have

Method	ME	RMSE	95th Percentile
Mean	15.777	92.379	171.742
Median	9.387	74.783	94.563
Minimum	-78.182	307.761	264.296
Proposed Method	**-1.504**	**30.251**	**59.305**

Table 1: Errors in ground height. All numbers in mm.

Step	Average Duration (ms)
Model construction	37.4
Histogram construction	16.2
New sample initialization	9.8
Resampling	86.7
Overheads	76.8
Total	226.9

Table 2: Average processing time for different steps in the ground height estimation. Overheads include interprocess communication from ROS, datatype conversions and coordinate transformations on point clouds.

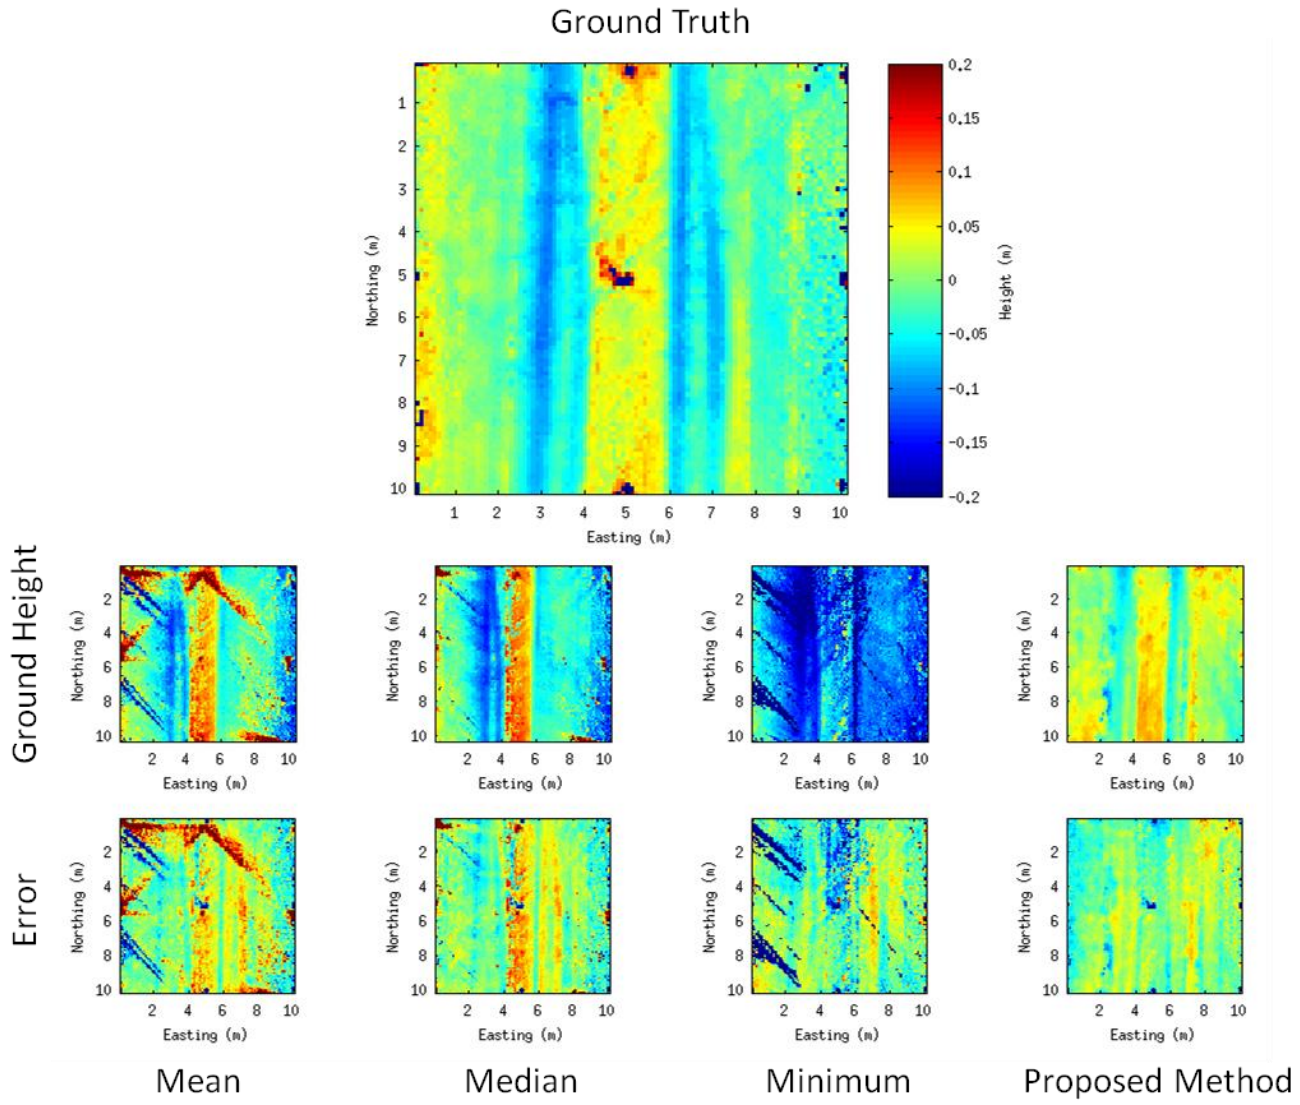

Figure 5: Predicted ground height using the proposed method, and a number of other naïve approaches. The colour indicates the height of the cell; blue is low and red is high. The large top figure shows the ground truth. The second row is the predicted ground height from each of the methods, and the third row is the error against the ground truth. The results are, left to right, mean, median, minimum, and the proposed method. The small circular regions in the ground truth figure correspond to the posts used for registration.

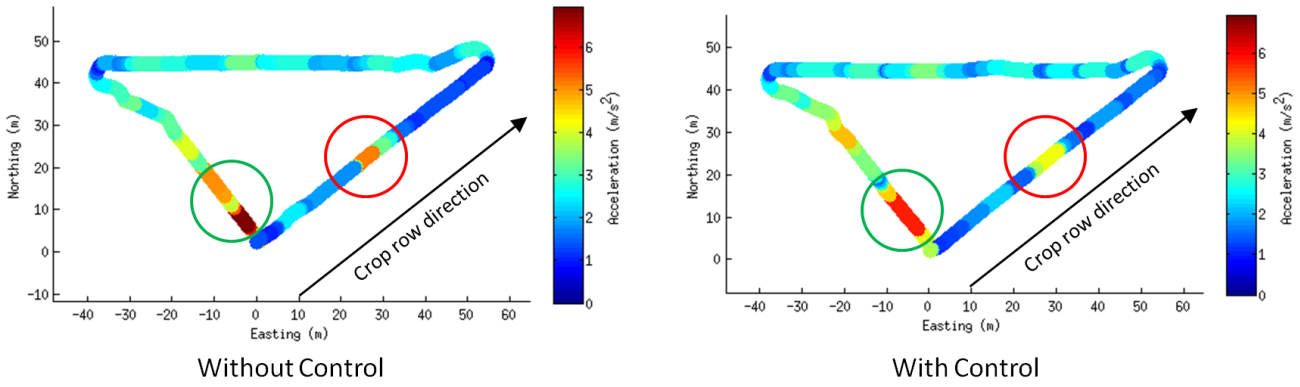

Figure 6: Vertical acceleration of the robot with and without the proposed speed control system. The red circled region indicates a deep ditch that the proposed method detected and slowed down the robot. . The green circled region indicates a region where the proposed method reduced the vertical acceleration of the robot for a hazard, although not less than the specified limit.

significantly lower error than the other methods in all respects, decreasing RMS error by 60%, the 95th percentile by 40%, and reducing the bias by 84%.

The limiting factor of the proposed method in this instance is that it oversmooths the higher frequency components of the ground. This could be mitigated by increasing the cell resolution or by increasing σ_g. This will be at the expense of increased processing time in the case of increasing the cell resolution, or increased estimation bias in the case of increasing σ_g.

Table 2 outlines the typical processing time for the algorithm on an Intel i7-2640M using only a single core. The total processing time per iteration allowed for an update frequency of 4.4Hz, which is approximately the same as stereo matching on the same machine.

5.2 Traversability controller

Figure 6 indicates the paths for each of the traverses and the vertical acceleration of the vehicle. The acceleration has been max-filtered for visibility in the figures. Overall the traversability controller successfully reduced the acceleration of the vehicle. The maximum acceleration was reduced by approximately 16% when applying the controller.

The significant hazard in this environment was a deep ditch (indicated on the diagram). The acceleration of the vehicle was reduced by 20% during the traversal of this region by slowing the robot.

There was a region when travelling across the crop rows where the acceleration of the vehicle wasn't limited as effectively by the algorithm. In this region the crop rows are much closer together, and as such are oversmoothed by the ground height estimator. As a result, this region appears to be much smoother. Additionally, the traversability metric only considers the acceleration of the wheel itself, assuming that it remains in constant contact with the ground. The acceleration of the vehicle is decoupled from the wheels by shock absorbers. This leads to inaccuracies between the predicted and measured accelerations, which is especially prevalent here as the vehicle is oscillating at close to resonance.

6 Conclusion

This paper has presented a novel algorithm for estimating the load-bearing surface in the presence of occlusions due to vegetation. It has been shown to significantly out-perform other strategies at estimating the true ground height in a challenging environment.

This method is applicable to many other strategies for traversability estimation, beyond just the one presented here. Many strategies for traversability estimation [Berczi *et al.*, 2015; Goldberg *et al.*, 2002; Ho *et al.*, 2013; Lacroix *et al.*, 2002; Tarokh and McDermott, 2005] rely on complete and accurate information of the location of the ground, and as such can be complemented by this strategy.

The presented method for traversability estimation has been shown to reduce the acceleration of the vehicle, however further work will be required to effectively limit the acceleration.

The algorithm requires a few environment-specific parameters in order to effectively estimate the ground height. Future work will investigate learning these parameters online from observations. Additionally, the model construction currently requires the robot to traverse a representative region of the environment in order to construct an accurate model. Future work will aim to mitigate this requirement.

Acknowledgements

This work was supported in part by the Australian Research Council Linkage Project LP110200375 "Robotics for Zero Tillage Agriculture" awarded to the Queensland University of Technology, SwarmFarm Robotics and the Australian Centre for Field Robotics.

References

[Angelova *et al.*, 2007] Anelia Angelova, Larry Matthies, Daniel Helmick, and Pietro Perona. Learning and Prediction of Slip From Visual Information. *Journal of Field Robotics (JFR)* 24(3): 205-231, 2007.

[Ball *et al.*, 2015] David Ball, Patrick Ross, Andrew English, Tim Patten, Ben Upcroft, Robert Fitch, Salah Sukkarieh, Gordon Wyeth, and Peter Corke. Robotics for Sustainable Broad-Acre Agriculture. *Field and Service Robotics (FSR)*, Springer International Publishing. 105: 439-453, 2015.

[Berczi *et al.*, 2015] Laszlo Berczi, Ingmar Posner, and Timothy Barfoot. Learning to Assess Terrain From Human Demonstration Using an Introspective Gaussian Process Classifier. *IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, USA, 2015.

[Geiger *et al.*, 2011] Andreas Geiger, Martin Roser, and Raquel Urtasun. Efficient Large-Scale Stereo Matching, Springer Berlin / Heidelberg. 6492: 25-38, 2011.

[Geman and Geman, 1984] Stuart Geman, and Donald Geman. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* PAMI-6(6): 721-741, 1984.

[Goldberg *et al.*, 2002] S. B. Goldberg, M. W. Maimone, and L. Matthies. Stereo Vision and Rover Navigation Software for Planetary Exploration. *IEEE Conference on Aerospace*, Big Sky, USA, 2002.

[Hadsell *et al.*, 2009] R. Hadsell, P. Sermanet, J. Ben, A. Erkan, M. Scoffier, K. Kavukcuoglu, U. Muller, and Y. LeCun. Learning Long-Range Vision for Autonomous Off-Road Driving. *Journal of Field Robotics (JFR)* 26(2): 120-144, 2009.

[Hastings, 1970] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57(1): 97-109, 1970.

[Ho *et al.*, 2013] K. Ho, T. Peynot, and S. Sukkarieh. Traversability Estimation for a Planetary Rover Via Experimental Kernel Learning in a Gaussian Process Framework. *IEEE International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany, 2013.

[Kim *et al.*, 2007] Dongshin Kim, Sang Min Oh, and J. M. Rehg. Traversability Classification for UGV Navigation: A Comparison of Patch and Superpixel Representations. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, San Diego, USA, 2007.

[Kim *et al.*, 2006] Dongshin Kim, Jie Sun, Sang Min Oh, J. M. Rehg, and A. F. Bobick. Traversability Classification

Using Unsupervised Online Visual Learning for Outdoor Robot Navigation. *IEEE International Conference on Robotics and Automation (ICRA)*, Orlando, USA, 2006.

[Kitt *et al.*, 2010] B. Kitt, A. Geiger, and H. Lategahn. Visual odometry based on stereo image sequences with RANSAC-based outlier rejection scheme. *IEEE Symposium on Intelligent Vehicles (IV)*, 2010.

[Lacroix *et al.*, 2002] Simon Lacroix, Anthony Mallet, David Bonnafous, Gerard Bauzil, Sara Fleury, Matthieu Herrb, and Raja Chatila. Autonomous Rover Navigation on Unknown Terrains: Functions and Integration. *International Journal of Robotics Research (IJRR)* 21(10-11): 917-942, 2002.

[Metropolis *et al.*, 1953] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics* 21(6): 1087-1092, 1953.

[Sofman *et al.*, 2006] Boris Sofman, Ellie Lin, J. Andrew Bagnell, John Cole, Nicolas Vandapel, and Anthony Stentz. Improving Robot Navigation Through Self-Supervised Online Learning. *Journal of Field Robotics (JFR)* 23(11-12): 1059-1075, 2006.

[Tarokh and McDermott, 2005] M. Tarokh, and G. J. McDermott. Kinematics Modeling and Analyses of Articulated Rovers. *Robotics, IEEE Transactions on* 21(4): 539-553, 2005.

[Vernaza *et al.*, 2008] P. Vernaza, B. Taskar, and D. D. Lee. Online, Self-Supervised Terrain Classification Via Discriminatively Trained Submodular Markov Random Fields. *IEEE International Conference on Robotics and Automation (ICRA)*, Pasadena, USA, 2008.

[Warren *et al.*, 2013] Michael Warren, David McKinnon, and Ben Upcroft. Online Calibration of Stereo Rigs for Long-Term Autonomy. *IEEE International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany, 2013.

[Wellington *et al.*, 2005] C. Wellington, A. Courville, and Anthony Stentz. Interacting Markov Random Fields for Simultaneous Terrain Modeling and Obstacle Detection. *Robotics: Science and Systems (RSS)*, Cambridge, USA, 2005.

[Wellington and Stentz, 2003] C. Wellington, and A. Stentz. Learning Predictions of the Load-Bearing Surface for Autonomous Rough-Terrain Navigation in Vegetation. *Field and Service Robotics (FSR)*, Springer Berlin Heidelberg. 24: 83-92, 2003.