18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Learned Ultra-Wideband RADAR Sensor Model for Augmented LIDAR-based Traversability Mapping in Vegetated Environments

Juhana Ahtiainen*, Thierry Peynot†, Jari Saarinen‡, Steven Scheding§, and Arto Visala*
* Department of Electrical Engineering and Automation, Aalto University, Finland. Email: {juhana.ahtiainen, arto.visala}@aalto.fi
† Science and Engineering Faculty, Queensland University of Technology, Australia. Email: t.peynot@qut.edu.au
‡ GIM Ltd., Finland. Email: jari.saarinen@gimltd.fi
§ Australian Centre for Field Robotics, University of Sydney, Australia. Email: s.scheding@acfr.usyd.edu.au

***Abstract*—In vegetated environments, reliable obstacle detection remains a challenge for state-of-the-art methods, which are usually based on geometrical representations of the environment built from LIDAR and/or visual data. In many cases, in practice field robots could safely traverse through vegetation, thereby avoiding costly detours. However, it is often mistakenly interpreted as an obstacle. Classifying vegetation is insufficient since there might be an obstacle hidden behind or within it. Some Ultra-wide band (UWB) radars can penetrate through vegetation to help distinguish actual obstacles from obstacle-free vegetation. However, these sensors provide noisy and low-accuracy data. Therefore, in this work we address the problem of reliable traversability estimation in vegetation by augmenting LIDAR-based traversability mapping with UWB radar data. A sensor model is learned from experimental data using a support vector machine to convert the radar data into occupancy probabilities. These are then fused with LIDAR-based traversability data. The resulting augmented traversability maps capture the fine resolution of LIDAR-based maps but clear safely traversable foliage from being interpreted as obstacle. We validate the approach experimentally using sensors mounted on two different mobile robots, navigating in two different environments.**

I. Introduction

Reliable obstacle detection is one of the most critical components of autonomous navigation in an a-priori unknown environment. Most existing obstacle detection systems rely on a geometric representation of the environment that is typically formed either by using a vision system or a LIDAR. However, for an unmanned ground vehicle (UGV) in vegetated environment this representation is not sufficient for effective navigation because vegetation is often mistakenly interpreted as an obstacle by perception systems based on these traditional sensors (visual cameras or LIDARs). In practice, traversing through sparse vegetation is often possible and preferable to avoid executing longer paths. In cases of densely vegetated environments it can even be the only acceptable option for the UGV. For example, in Fig. 1 the UGV, relying on a state-of-the-art terrain traversability mapping system based on LIDAR data, tends to consider the grass as an obstacle and would not be able to traverse the surrounding terrain.

Fig. 1. UGV in a densely vegetated environment.

Prior studies have approached this problem by classifying vegetation to distinguish it from other types of obstacles. For example, in [1] a multispectral camera is utilised for detecting vegetation based on near-infrared (NIR) light reflectance properties. In [2] and [3] a 3D LIDAR is used to classify grass from other obstacles based on statistical analysis of the 3D data points. However, solid obstacles might be hidden behind the vegetation (e.g. a large rock), which would pose a great risk to the robot if it were to decide to traverse through this vegetated area based on the result of vegetation classification only. Some ultra-wideband (UWB) radars, which operate at much lower frequencies of the electromagnetic spectrum, are able to see through a certain amount of vegetation [4]. However, these sensors emit large beams and provide data with high noise and low accuracy. Therefore, UWB radars alone are insufficient to provide accurate terrain traversability maps.

In prior work, the authors showed that by combining data acquired by a UWB radar and a LIDAR it was possible to generate *augmented* traversability maps that allow for reliable navigation of UGVs in vegetated environments [5]. LIDAR-based traversability maps were augmented with radar measurements such that areas with obstacle-free foliage (i.e. areas of vegetation that the UGV could safely drive through) could be cleared from the map. A sensor model was developed to convert the radar measurements into occupancy probabilities, which were fused with probabilities of occupancy computed

from LIDAR measurements. The radar measurements were used only to update the cells seen as obstacles by the LIDAR, which reduced the effect of the UWB radar noise. However, this sensor model was partly engineered manually, and had to be tuned by an operator. This is a tedious operation that may need to be repeated for each mounting configuration of the radar and/or type of environment. Therefore, in this paper we propose to learn the UWB radar sensor model from experience, using a support vector machine (SVM) [6]. We provide an experimental validation of the proposed method using two different UGV platforms in two distinct environments. We show that thanks to the learned sensor model the new augmented traversability maps clearly outperform the state of the art.

The paper is organized as follows. Section II reviews the relevant prior work. Section III describes the proposed approach in detail. Section IV introduces the experimental system and discusses implementation details, then Sec. V presents the experimental results and compares them to previous work. Finally, Sec. VI concludes the paper and proposes some future work.

II. Related Work

The most common representation of the environments for obstacle detection of mobile robots is probabilistic occupancy grid maps [7], which typically divide the environment into equal grid cells whose values are probabilities of being occupied (by an obstacle) or not. To populate these maps, LIDARs are the most commonly used sensors, however, any other range sensor may be used instead [8], or in combination with LIDARs. For example, millimetre-wave (MMW) radars were used in [9] and [10] to construct probabilistic occupancy grids in environments that are challenging for laser or cameras, such as in the presence of heavy dust, where LIDARs tend to fail to provide reliable information beyond (or through) dust clouds [11], [12].

Traversability maps quantify the difficulty a robot would encounter when traversing through a particular region. They are typically platform-dependent [13]. In [14] an elevation map, i.e. a grid-based $2.5D$ spatial representation where each cell stores the height of the cell [15], was transformed into a traversability map where each grid cell retains a traversability index. This index was calculated using the terrain slope and roughness. In [16] a machine learning method was exploited to learn the traversability of a road ahead using data from LIDAR, camera, and inertial measurement unit (IMU). All these methods rely only on range data from sensors that are not capable of consistently penetrating foliage. Therefore, they tend to interpret foliage as difficult areas for a UGV, if not obstacles.

UWB radars have only been utilized in few studies for detecting obstacles in vegetated areas. In [17] an impulse radar operating at $2.2\,GHz$ was used to detect a tree trunk behind $2.5\,m$ of branches and foliage. A ground penetrating radar was used to detect obstacles within vegetation as well as underground in [18]. To study the possibility of detecting obstacles through vegetation, a custom UWB radar array was built and tested in [19] and [20]. However, these studies concentrate only on UWB-radar-based obstacle detection and do not take the terrain traversability into consideration. In iRobot's DareDevil project [21] the authors used a UWB radar in parallel with a LIDAR for all-weather operations of a UGV, and also discussed the possibility of achieving obstacle detection within vegetation by comparing the output of the two sensors. However, this was not actually demonstrated.

III. Approach

We propose a method for learning the sensor model for a UWB radar from data using a SVM. The learned sensor model is used to convert the UWB radar measurements to occupancy values that are fused with the LIDAR data to augment LIDAR-based traversability maps in vegetated environments. This section presents the details of the learning process as well as how the learned sensor model is applied.

A. Target Detection with UWB Radars

UWB radar systems transmit signals across a much wider range of frequencies than conventional radar systems. The most common technique for generating UWB signals is to transmit pulses with very short durations. These pulsed UWB radars return a vector of power measurements originating from the radar cross-sections (RCS) of the targets within the radar field of view (FOV) [9]. The elements of this vector are referred to as range bins. Typically UWB radars operate at low frequencies and are therefore capable of penetrating some amount of soft material, e.g., vegetation. On the other hand, the low frequency results in a wide beam [4], which makes it challenging to accurately localise detected targets, and to estimate their actual dimensions. Furthermore, the UWB radar return vectors can be very noisy.

Traditionally, the target detection from radar data is based on comparing the radar returns to threshold values. The popular constant false alarm rate (CFAR) method [22] does not perform well in environments with frequent obstacles or with radars that return short measurement vector. The thresholds presented in [10] cannot be adapted to the case of UWB radars due to different noise characteristics [4]. The sensor model introduced by the authors in [5] was shown to be appropriate for a UWB radar. However, the model was partly engineered by hand which lacks adaptability properties and may lead to sub-optimal performance. Therefore, we propose that the sensor model is learned from experience, using real data, before the radar data are combined with the LIDAR measurements.

B. Learning the Sensor Model

The UWB radar sensor model is learned based on the synchronized and localized radar and LIDAR data from controlled environment using the following five steps.

1) Generate LIDAR-based Traversability Map: The UGV, equipped with a LIDAR pointed towards the ground, a UWB radar, and a localisation system, is driven around a controlled

environment to collect range data with corresponding localisation. We assume the relative configuration of the sensors on the platform is know. This can be estimated by a calibration method such as in [23]. A 2D traversability grid map, T_m, is then generated from the LIDAR measurements. The radar data needs to be recorded at the same time. Each cell on this map contains a traversability value, τ, that quantifies the difficulty for a robot to travel through that area. The traversability map is normalised such that $\tau \in [0,1]$.

2) Label Teaching Examples: Traversable and untraversable areas are then hand-labelled on the grid map, which is straightforward since the environment is controlled and safely traversable areas are known. There is no need to label the whole map but only representative examples from both classes. That is, to ensure the quality of the training data, only the areas whose labels are definitely known should be labelled. In addition, labelling multiple smaller areas enables a more even distribution of the training data. This is discussed further in the next section.

3) Extract Features: Features are extracted based on the labelled cells of the grid, the radar return vectors, the radar FOV, and the synchronized radar poses. The features we use are the angle from the radar to a labelled cell, α, the corresponding range bin index, k, and the measured intensity at this range bin $I(k)$. The radar data are replayed with respect to the synchronized radar pose and the labelled cells within the effective radar FOV are taken into account. An illustration of this process can be seen in Fig. 2.

Typically there are more than one labelled cell corresponding to a range bin simultaneously. Therefore, only the cells with the smallest α are saved for each range bin at every time step since it is assumed that the strongest reflection is received from the middle of the FOV. Traversable cells are only saved if there are no labelled obstacle cell at the corresponding range bin. For example, in Fig. 2, only O_1, T_2, and O_2 are saved for range bins $k-1$, k, and $k+1$ respectively. In case the whole map would be labelled, only teaching examples with very small α would be recorded. Therefore, in order to get more even distribution of α, it is important to label multiple smaller regions rather than one large one. Nonetheless, the different labelled classes should have clear margin between them to avoid confusion.

Since the sensors are calibrated with respect to the same frame, we have the timestamped radar poses in the same coordinate frame as the labelled LIDAR-based grid map. Therefore, it is straightforward to calculate the α and the distance d. The range bin number is calculated by $k = \lceil d/\Delta d \rceil$, where Δd is the range resolution. The corresponding intensities are fetched from the radar return vector. Note that we discard all the labelled cells that are not within the effective radar detection range $[d_{min}, d_{max}]$. This is determined experimentally based on the lack of reliable detection at ranges beyond d_{max} and the excess of noise in the data closer than d_{min}.

4) Preprocess Data: The extracted data need some preprosessing before they can be used for teaching the model. Namely, the number of negative and positive samples might

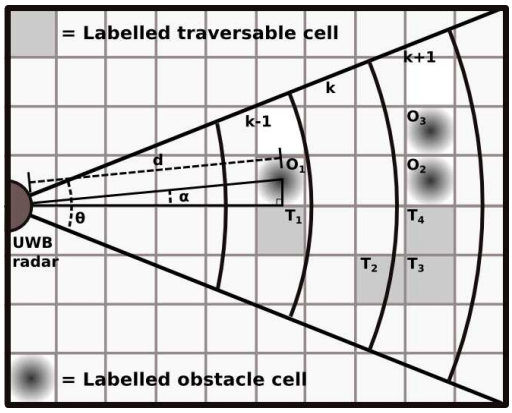

(a) Radar FOV

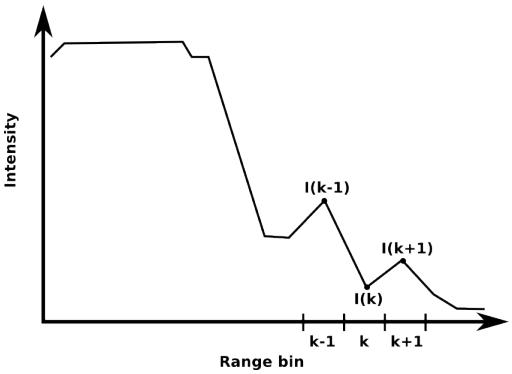

(b) Corresponding radar return

Fig. 2. (a) An illustration of the radar FOV while extracting the features with labelled cells on the FOV ($T_1 - T_4$ and $O_1 - O_3$). Three range bins are depicted in the figure ($k-1$, k, and $k+1$). θ is the radar beamwidth, α is the angle from the radar to a labelled cell, and d is the distance to a labelled cell. (b) Corresponding radar return vector. $I(k-1)$, $I(k)$, and $I(k+1)$ are the intensities of the range bins.

be significantly unbalanced. Therefore, the larger group is randomly downsampled such that the number of examples from both classes are roughly equal for each range bin. This ensures that no problems arise from unbalanced data. In addition, we scale the features within $[0,1]$ and save the original minimum and maximum values for each k such that we can scale the measurements similarly when applying the learned sensor model.

5) Train the Models: The preprocessed data are used to train a model for converting the radar data to measurement likelihoods. The well-known SVM framework was chosen as the learning method because it has been shown that SVMs consistently perform very well in similar two-class real-world problems [24]. We use the C-support vector classification (C-SVC) [6] with the probability estimates extension [25] to estimate a probability that an object in the radar FOV is indeed a real obstacle. We train a separate model for each range bin since the radar-return-vector noise floor differs significantly from one range bin to another. The separate models are simpler than one complex model, which makes the prediction step

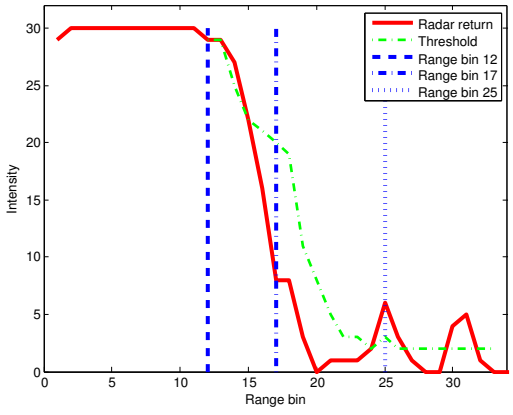

(a) Radar return

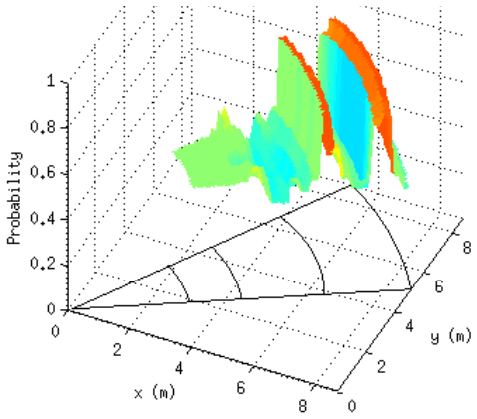

(b) Corresponding occupancy probabilities

Fig. 3. (a) A radar return with two targets in the FOV (at a range of 7.5m and 9m, respectively) drawn in red. The learned threshold (assuming $\alpha = 0$) is drawn in green. (b) Corresponding occupancy probabilities calculated with the learned sensor model.

much faster. We use Radial Basis Function (RBF) kernel since it can handle non-linear relations in the data and tends to perform well in cases where there are only few features [26].

The free parameters for the C-SVC with RBF kernel are C, which controls the cost of misclassification on the training data, and the kernel parameter γ, which controls how far the influence of a single training example reaches. We used grid search guided by cross-validation to assist in the selection of these parameters. Note that since the radar data are very noisy, it is important not to overfit the model and aim for smooth decision surfaces. An example of the radar return vector and the learned sensor model can be seen in Fig. 3. The green detection threshold in Fig. 3(a) was computed assuming that measurements originate from the middle of the FOV and it illustrates the point where the predicted class changes.

C. Apply the Model

Applying the learned sensor model is done in three steps. Firstly, T_m is generated from the LIDAR data. Secondly, the sensor model is applied for each untraversable cell on T_m within the radar FOV. Thirdly, the radar detections are combined with the prior data.

1) Generate LIDAR-based Traversability Map: First, as in the learning phase, we need to generate T_m from the LIDAR data. However, contrary to the learning phase, a local traversability map of the robot's surroundings is sufficient as long as it covers the radar FOV. Constantly updating a local T_m enables real-time processing of the proposed algorithm.

2) Apply the Sensor Model: After each new radar measurement vector, all the untraversable cells in T_m within the radar FOV are processed with the sensor model. That is, the features α, k, and $I(k)$ are extracted from the data and scaled, the correct model is selected based on the value k (the range bin number), and the model is applied using the feature vector. However, there are often multiple target cells that correspond to a particular range bin similarly to Fig. 2. This is problematic since the radar only returns the strongest reflection per each range bin. Therefore, we assume that the $I(k)$ originates from the cell with the smallest α. To take this into account, we apply an angle scaler G_θ that models the beam pattern with an inverse parabola [9] and scales the probabilities with respect to the cell with the smallest α. That is, G_θ is

$$G_\theta(\alpha, \alpha_{min}) = \frac{\theta^2 - 2\alpha^2}{\theta^2 - 2\alpha_{min}^2}\,, \tag{1}$$

where α and α_{min} are the angle from the beam axis to the centre of the current cell and the central cell respectively, and θ is the beamwidth of the radar. The measurement likelihood, $P(z|y)$, is calculated with

$$P(z|y) = G_\theta(\alpha, \alpha_{max}) \times P_{pred}(y)\,, \tag{2}$$

where the class labels $y \in \{0, 1\}$ correspond to traversable and untraversable cells respectively. The $P_{pred}(y)$ is the predicted class probability and the z is the radar return vector. The measurement likelihood is computed only for the predicted y and the result is bounded such that $0.5 \leq P(z|y) \leq 0.8$. The upper bound controls the conservativity of the clearing strategy: the lower the value, the more measurements are required to clear cells initially found as obstacles in the LIDAR-based map.

3) Update the Augmented Traversability Map: The corresponding cells are updated to the augmented traversability map, T_{ma}, by using Bayes' formula, assuming static world and conditional independence, a standard assumption in occupancy mapping applications [7]. The T_{ma} is initialized with the values from T_m but only the cells seen untraversable by the LIDAR are updated. This ensures that the resulting T_{ma} retains the high resolution of the T_m and the noisy radar measurements are only used for additional information on the obstacles. In case the radar observations do not indicate that a cell should be cleared (i.e., the predicted probabilities are not consistently low enough), the cell remains untraversable.

IV. System Description

This section presents the experimental system used for validation and discusses the implementation of the proposed approach on two UGVs.

A. UWB Radar

The UWB radar used in this study is a Radar Developer's Kit Lite (RaDeKL) by Multi Spectral Solutions Inc (MSSI). The most relevant radar performance characteristics are summarised in Table I. Please refer to [5] for an experimental evaluation of the ability of this sensor to penetrate through various depths of vegetation.

TABLE I
TECHNICAL SPECIFICATIONS OF RADEKL UWB RADAR

RF Characteristics	
Centre Frequency	6.35 GHz
Bandwidth	400 Mhz (-3 dB)
Peak Power	50 mW EIRP
Antenna gain	12 dBi w/4x4 Array
Antenna FOV	40 deg AZ x 40 deg EL
System Performance	
Range Extent	256 range bins w/variable offsets
Range resolution (Δd)	30 cm
Effective min range (d_{min})	3.6 m
Effective max range (d_{max})	10.0 m

B. Measurement Platforms

Data from two different measurement platforms were used in this study to validate the proposed approach. The measurement platforms, Shrimp and Argo, can be seen in Fig. 4. Shrimp is based on Segway's Robotic Mobile Platform RMP-400 and Argo is an 8 wheel skid-steering vehicle. Both platforms are equipped with a Novatel SPAN System (Synchronized Position Attitude and Navigation) with a Honeywell IMU positioning system, which typically provides $2\,cm$-accuracy localisation estimates. On both platforms the UWB radar was mounted at the front of both platforms, slightly pointing down (tilt angle of about $2\,deg.$). A SICK LMS 291 LIDARs (indicated in Fig. 4), mounted with a tilt angle of about $8\,deg.$, was used to generate the LIDAR-based traversability maps by scanning the terrain with the vehicle.

C. Implementation

In this paper, the traversability map, T_m, was computed using the method in [14]. First, an elevation map was calculated from the LIDAR returns and the synchronized pose data. The traversability index, τ, was then calculated for each cell using the slope and roughness of the terrain.

The resulting T_m was scaled by dividing the τ of each cell by a platform-specific untraversability constant, $T_i = 40$, which was determined experimentally. Every cell with a τ above T_i was considered untraversable and the values below were scaled between 0 and 1. However, all traversability representations are applicable as long as the traversability is scaled accordingly.

We used the popular LIBSVM library [27] to learn the sensor model. The range bin models were divided into three groups with similar noise characteristics based on radar return data analysis and the grid search results. The selected hyperparameters and the mean cross-validation accuracies are presented in Table II. The γ parameter of the C-SVC method

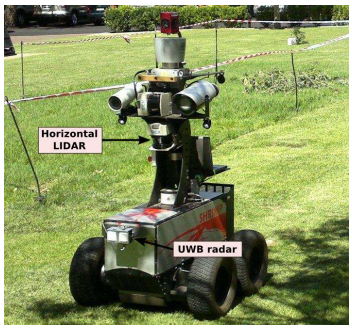

(a) Shrimp

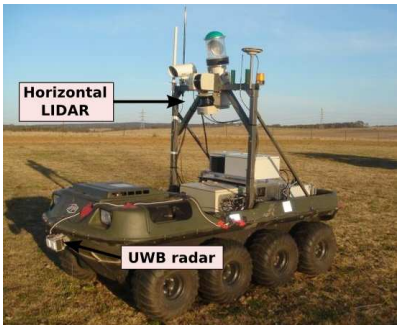

(b) Argo

Fig. 4. The measurement platforms used for validation. The exteroceptive sensors used in this study are marked on the pictures.

TABLE II
HYPERPARAMETERS OF THE SENSOR MODELS (*see Sec. III-B5)*

Range bin	γ	C	Accuracy
12-16	0.50	0.50	61%
17-24	0.25	1.00	73%
25-33	0.50	8.00	84%

(see Sec. III-B5) is low for all the range bins since the underlying non-linear model was expected to be simple and remain similar throughout the data. However, the range bins closer to the radar have significantly higher noise floor. Therefore, separating the targets from the noise is more challenging and the decision surface was kept smooth (i.e., C is low) to avoid overfitting.

V. EXPERIMENTAL RESULTS

To validate the proposed approach, we use the experimental data presented in [5]. Three sets of field trials were conducted. Trials 1 and 2 were conducted in a controlled environment on a relatively flat lawn and Trial 3 in a rural environment with numerous grass tussocks on the test area (see Fig. 5). We used the data from Trial 1 for learning the sensor model and the data from Trials 2 and 3 to validate the proposed approach. The results are also compared against the method in [5].

In Trial 2, a total of 6 different experiments were performed. In experiment a, the lawn was clear of obstacles and three obstacles (i.e., a stone and columns of 2 and 3 bricks) were

(a) Test site of Trials 1 and 2 (b) Test site of Trial 3

Fig. 5. Partial views of the test sites.

TABLE III
CLEARED CELLS FROM TRIAL 2

	Learned sensor model		Sensor model in [5]	
	TNR	*FNR*	*TNR*	*FNR*
a) Empty	N/A	**6.67%**	N/A	6.81%
b) Obstacles	N/A	**12.68%**	N/A	15.42%
c) 1 layer	**92.40%**	9.01%	88.98%	**8.59%**
d) 2 layers	**92.86%**	**14.21%**	88.49%	17.60%
e) 3 layers	**83.96%**	**12.25%**	75.00%	13.87%
f) 4 layers	**83.09%**	14.85%	65.69%	**10.37%**

added for the experiment b. Then, branches of Eucalyptus tree were added layer by layer in front of the obstacles as well as on two clear spots on the lawn for experiments c, d, e, and f. One layer was added before each experiment, every layer adding approximately $10\,cm$ of depth to the foliage. See the summary of experiments in this trial in Table III.

An example of the resulting T_m from Trial $2c$ (with one layer of foliage in front of obstacles) is shown in Fig. 6(a). The locations of the obstacles are indicated in the figure. The corresponding augmented traversability map, T_{ma}, is shown in Fig. 6(c).

Fig. 6(b) indicates how long the map cells that were updated using the radar data spent in the radar FOV during the experiment. Note that due to the noise and low resolution of the radar data, and to use a sufficiently conservative strategy, the decision of clearing a grid cell from being considered as an obstacle was only taken for cells that have been in the radar FOV for more than $2\,s$.

Table III summarises the results from Trial 2 using our approach, compared with the hand-tuned sensor model in [5]. Since our goal is to clear some areas falsely considered as obstacled in the LIDAR-based traversability map T_m, in this analysis we are only interested in the false negative rate (FNR) and the true negative rate (TNR). TNR is the proportion of foliage cells that were correctly cleared of being an obstacle, i.e. obstacle-free foliage cells, among all obstacle-free foliage cells. FNR is the proportion of obstacle cells that were falsely cleared, among all obstacle cells. These values are calculated based on hand labelling foliage and obstacles on the T_m using information of the obstacle locations and size without foliage. In the table, the best results for each experiment (a-f) (i.e. highest TNR and lowest FNR) are shown in bold font.

It is clear from Table III that our approach outperforms the model in [5] in all the experiments. The TNR values are significantly higher with the learned sensor model, i.e., more obstacle-free vegetation is correctly cleared. The TNR values are around 90% in the experiments with less than three layers of foliage but start to drop as the vegetation increases since the radar is no longer able to distinguish between foliage and obstacles.

The FNR values are around 10% with both approaches. Overall, there are three reasons that affect the FNR values. Firstly, the radar is not able to capture the true dimensions of the objects due to large range resolution, which may result in falsely clearing some of the obstacle cells. Secondly, the annotation is based on traversability map that typically exaggerates the obstacle dimensions, i.e., some of the perimeter cells of labelled obstacles might be obstacle-free in reality. Finally, the RCS of some of the obstacles is too small for the radar to detect them reliably and in some experiments they are falsely cleared. For example, in experiment f with 4 layers of vegetation (see Table III) the column of two bricks was falsely cleared with the learned approach, which leads to a significant increase in the FNR due to the low number of labelled cells.

Trial 3 was conducted in a rural environment with multiple grass tufts and three brick piles (heights: 2, 3, and 4 bricks) on the test area. In addition, there was also a small ditch and a car in the test area. Maps of Trial 3 can be seen in Fig. 7. The LIDAR-based traversability map in Fig. 7(a) shows that when using only LIDAR data the terrain appears very challenging and if the UGV was located in Position A it would not find a safe path to autonomously reach Position B, even though most of the field is actually only covered with grass, with only occasional and isolated actual obstacles. However, augmenting the traversability map with UWB radar data allowed the system to clear large areas that have been observed by the radar (Fig. 7(b)) and considered as obstacle-free despite the presence of tall grass (Fig. 7c-d). Both approaches performed reasonably well in this test, however, the new approach with learned sensor model was able to clear significantly more obstacle-free cells, while still detecting all actual obstacles. The TNR and FNR could not be computed for Trial 3 since no ground truth was available for these tests.

VI. CONCLUSION

In this paper, we proposed a method for learning the UWB radar sensor model used to augment LIDAR-based traversability maps in vegetated environments. This is especially important in densely vegetated environments where it may be impossible for a UGV to operate without sensors that are able to penetrate some amount of vegetation. A SVM with probability extension was utilized to learn the sensor model that converts the radar measurement vector into occupancy values of individual cells that can be fused with the LIDAR measurements. We showed that the learned sensor model allowed us to outperform the method introduced in our prior work, especially increasing significantly the True Negative Rate value, i.e. clearing obstacle-free foliage areas

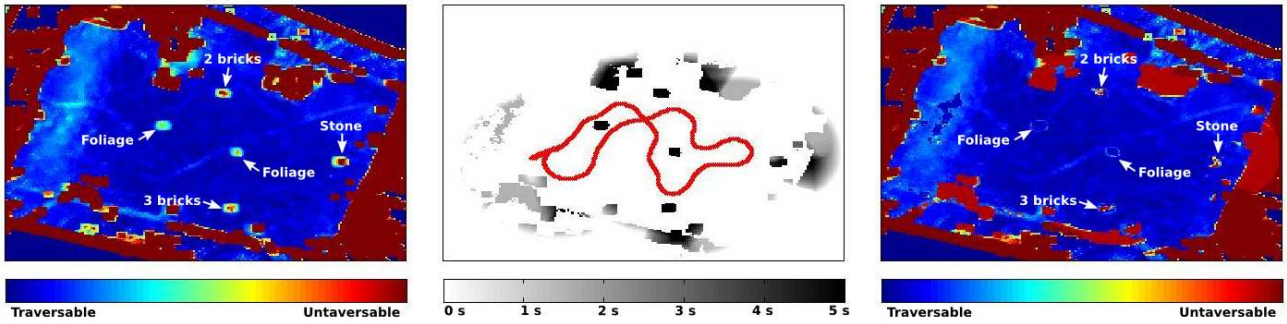

(a) Traversability map (T_m) (b) Updated cells radar FOV (c) Augmented map (T_{ma})

Fig. 6. Maps from Trial 2*c*. (a) shows the LIDAR traversability map, coloured by traversability value (red means obstacle); (b) show the cells that were observed by the radar, with an intensity of grey proportional to the time spent by the cell in the radar FOV (darker means longer time); (c) show the augmented occupancy map, coloured by probability values, (blue for 0, red for 1). The size of the test area is around $30 \times 30 m^2$.

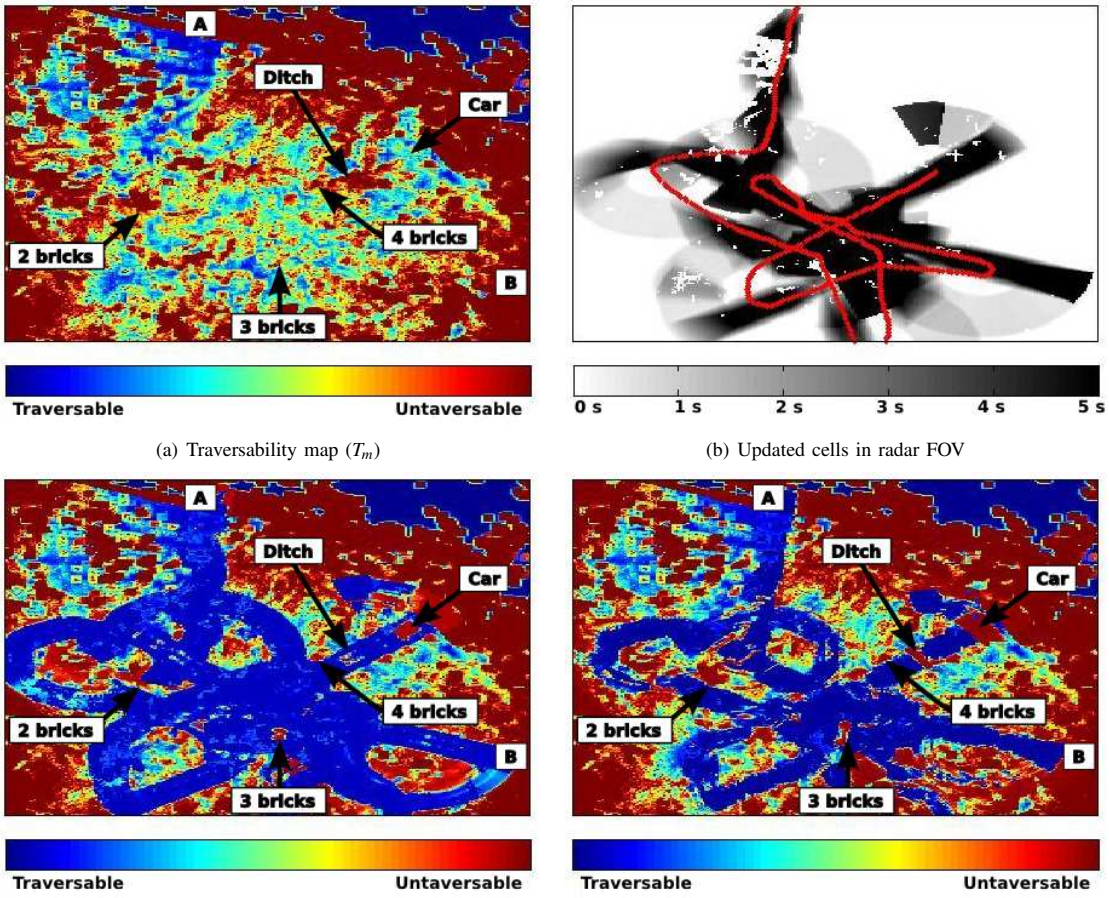

(a) Traversability map (T_m) (b) Updated cells in radar FOV

(c) T_{ma} using learned sensor model (d) T_{ma} using method in [5]

Fig. 7. Trial 3: (a) shows the LIDAR traversability map, coloured by traversability value (red means obstacle); (b) shows the cells that were observed by the radar, with an intensity of grey proportional to the time spent by the cell in the radar FOV (darker means longer time); (d) shows the augmented occupancy map using the learned sensor model, coloured by the probability value (blue for 0, red for 1). (d) shows the augmented occupancy map using the model in [5], coloured by the probability values. The size of the test area is around $50 \times 50 m^2$.

that were considered as obstacles by state-of-the-art LIDAR-based traversability mapping.

However, despite the clear improvement of performance the radar data remain too noisy to reliably distinguish targets with small radar cross-sections, which may lead to ignoring small difficulties in the terrain detected by the LIDAR. Even though these targets are typically not a significant threat to the integrity of the robot, it is preferable to detect them to anticipate any difficulty (e.g. to reduce the speed of the platform). Therefore, the learned sensor model is trained to clear vegetation aggressively while still detecting the obstacles in the experiments. More aggressive sensor model would lead to falsely clearing more obstacles.

The learned sensor model is applicable to various environments with different types of vegetation. However, the amount of cleared obstacle-free vegetation may depend on the type of the vegetation. For example, wet grass or vegetation with high concentration of branches will result in stronger backscattering of the radar signal, i.e., less obstacle-free vegetation can be cleared.

To improve performance, in future work we consider mounting the radar on a scanning mechanism (e.g. a pan-tilt unit) or using an overlapping array of radars, which would enhance the accuracy of the collected radar data, as well as that of the augmented traversability map. We will also consider implementing a real-time version of the algorithm. In this paper the results were obtained off-line, however, no heavy computations are involved in the process, therefore, using a 3D LIDAR should allow us to update a local traversability map continuously.

ACKNOWLEDGMENT

This work was partly supported by the Australian Centre for Field Robotics and the Academy of Finland. Their support is gratefully acknowledged. The authors thank Doug Warwick for providing vegetation, Peter Morton and Zhe Xu for their help with the platforms, Mark Calleija for mounting the radar, and Yannick Godin for photographs and data logging.

REFERENCES

[1] D. Bradley, S. Thayer, A. Stentz, and P. Rander, "Vegetation detection for mobile robot navigation," Robotics Institute, Carnegie Mellon University, Tech. Rep., 2004.

[2] A. Castano and L. Matthies, "Foliage discrimination using a rotating ladar," in *Proc. of IEEE Int. Conf. on Robotics and Automation*, Taipei, Taiwan, 2003.

[3] J. Macedo, R. Manduchi, and L. Matthies, "Laser-based discrimination of grass from obstacles for autonomous navigation," in *Proc. of Int. Sym. on Experimental Robotics*, Honolulu, HI, 2000.

[4] J. Taylor, *Introduction to ultra-wideband systems*. Florida: CRC Press LLC, 1995.

[5] J. Ahtiainen, T. Peynot, J. Saarinen, and S. Scheding, "Augmenting traversability maps with ultra-wideband radar to enhance obstacle detection in vegetated environments," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Tokyo, Japan, 2013.

[6] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, 1995.

[7] H. Moravec, "Sensor fusion in certainty grids for mobile robots," *AI Magazine*, vol. 9, no. 2, 1988.

[8] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[9] A. Fossels, "Scene modelling from motion-free radar sensing," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, 2002.

[10] B. Clarke, S. Worrall, G. Brooker, and E. Nebot, "Sensor modelling for radar-based occupancy mapping," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Vilamoura, Portugal, 2012.

[11] T. Peynot and A. Kassir, "Laser-camera data discrepancies and reliable perception in outdoor robotics," in *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Taipei, Taiwan, 2010.

[12] M. P. Gerardo-Castro and T. Peynot, "Laser-to-radar sensing redundancy for resilient perception in adverse environmental conditions," in *Proc. of ARAA Australasian Conf. on Robotics and Automation*, Wellington, New Zealand, 2012.

[13] V. Molino, R. Madhavan, E. Messina, A. Downs, S. Balakirsky, and A. Jacoff, "Traversability metrics for rough terrain applied to repeatable test methods," in *Proc. of IEEE/RSJ Int. Conf. on Robotics and Automation*, Rome, Italy, 2007.

[14] C. Ye and J. Borenstein, "T-transformation: a new traversability analysis method for terrain navigation," in *Proc. of the SPIE: Unmanned Ground Vehicle Technology VI*, Orlando FL, 2004.

[15] B. Siciliano and O. Khatib, *Springer Handbook of Robotics*. Springer, 2008.

[16] M. Montemerlo, S. Thrun, H. Dahlkamp, and D. Stavens, "Winning the DARPA grand challenge with an AI robot," in *Proc. of AAAI National Conf. on Artificial Intelligence*, Boston, MA, 2006.

[17] L. Matthies, C. Bergh, A. Castano, J. Macedo, and R. Manduchi, "Obstacle detection in foliage with ladar and radar," in *Proc. of Int. Sym. Laboratory Research*, Sienna, Italy, 2003.

[18] A. Boryssenko, "Preventing damage by hidden objects in vegetation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 21, no. 12, 2006.

[19] S. Nam, S. Sun, and G. C. Park, "Frequency and polarization characteristics in vegetation for the ground based penetrating radar," in *Proc. of Int. Asia-Pacific Conf. on Synthetic Aperture Radar*, Seoul, Korea, 2011.

[20] S. G. Sun, B. Cho, G. C. Park, Y. S. Kang, and S. H. Han, "UWB forward imaging radar for an unmanned ground vehicle," in *Proc. of Int. Asia-Pacific Conf. on Synthetic Aperture Radar*, Seoul, Korea, 2011.

[21] B. Yamauchi, "Fusing ultra-wideband radar and LIDAR for small UGV navigation in all-weather conditions," in *Proc. of SPIE: Unmanned Systems Technology XII*, Orlando, FL, 2010.

[22] H. Rohling, "Radar CFAR thresholding in clutter and multiple target situations," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 19, no. 4, 1983.

[23] J. P. Underwood, A. Hill, T. Peynot, and S. J. Scheding, "Error modeling and calibration of exteroceptive sensors for accurate mapping applications," *J. Field Robot.*, vol. 27, no. 1, 2010.

[24] T. Van Gestel, J. A. K. Suykens, B. Baesens, S. Viaene, J. Vanthienen, G. Dedene, B. De Moor, and J. Vandewalle, "Benchmarking least squares support vector machine classifiers," *Mach. Learn.*, vol. 54, no. 1, 2004.

[25] T.-F. Wu, C.-J. Lin, and R. C. Weng, "Probability estimates for multi-class classification by pairwise coupling," *J. Mach. Learn. Res.*, vol. 5, 2004.

[26] C.-W. Hsu, C.-C. Chang, and C.-K. Lin, "A practical guide to support vector classification," Department of Computer Science, National Taiwan University, Tech. Rep., 2010.

[27] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, 2011.