Stochastic Classification of Urban Terrain for Smart Wheelchair Navigation

Corey Montella, Michael Pollock, Dylan Schwesinger and John R. Spletzer
Department of Computer Science and Engineering
Lehigh University

***Abstract*—In this paper, we investigate the potential for enhancing the perception capabilities of smart wheelchair system (SWS) in an urban environment through intensity-based terrain classification. A 3D camera system was employed that provided both range and intensity (remission) measurements to objects in the scene. The feature space of the sensor measurements was decomposed into range ρ, bearing angle ϕ, and intensity value γ. We then constructed stochastic models by discretizing the feature space into cells, each of which corresponded to a probability mass function (PMF) denoting the probability a point in the feature space belonged to one of four classes: grass, sidewalk, street (blacktop), or unknown. The stochastic models were generated *a priori* from literally millions of labeled point measurements of the different terrain classes. Bayesian filtering techniques were then used to make a terrain class assignment for ground plane cells in the local map over time. We illustrate the utility of the approach by demonstrating autonomous sidewalk following of the SWS over hundreds of meters solely through terrain class assignment. Extensions such as enhanced ground plane segmentation are also discussed.**

I. Introduction

Smart wheelchair systems (SWS) have been an active research area for nearly three decades. While the breadth of research is too great to summarize here, a survey as of August 2005 by Simpson can be found at [1]. More recent projects of note include the MIT Intelligent Wheelchair Project [2], the goal of which is to develop a voice-commanded autonomous wheelchair intended for use in indoor environments. The Home, Lift, Position, and Rehabilitation (HLPR) Chair [3] being developed by NIST is a special-purpose device to provide independent patient mobility for indoor tasks, such as moving to and placing a person on a chair or bed. The Personal Mobility and Manipulation Appliance (PerMMA) [4] is being developed at the University of Pittsburgh and Carnegie Mellon University with the objective of combining manipulation and mobility assistance in support of complete independence for its users. Despite these efforts, commercial SWS solutions remain elusive. One exception is our work on the Automated Transport and Retrieval System (ATRS) [5]. Through the use of robotics and automation technologies, ATRS improves automobile access for power wheelchair users by eliminating the need for an attendant to stow and retrieve the wheelchair. The system received U.S. Food and Drug Administration (FDA) approval in 2008, and is currently available commercially [6].

The emphasis of our current work is autonomous navigation of SWS in urban environments. In [7], we showed that using a map-based localization approach, our SWS could reliably navigate from Lehigh's Packard Laboratory to the university bookstore, a distance of approximately 1 km round trip. During these experiments, we also noted environmental features that could lead to system failure under certain conditions. For example, a sidewalk bordered by grass or mulch beds would often be entirely classified as sidewalk based upon ground plane geometry. A similar limitation of ground plane segmentation approaches was also seen in sidewalk areas with low-lying curbs. At times, transitions between the street and sidewalk could not be detected.

The take-home message from these scenarios is that *geometry alone is insufficient for robust navigation in urban environments.* To be certain, we are not alone in reaching this conclusion. In a study on accessibility in urban settings, Meyers *et al*. found that the most common environmental barriers to wheelchair navigation are steep ramps, sidewalk obstructions, rough terrain, total travel distance, and lack of curb cut outs [8]. Therefore, an appropriate map to facilitate sidewalk-level navigation must include not only standard geometric information, but also semantic labels delineating features like terrain type, curbs and cutouts, and sidewalk grade. To this end, in this work we investigated the potential for augmenting the navigation performance of a SWS through intensity-based terrain classification. A 3D camera system was employed in conjunction with Bayesian filtering techniques to apply a stochastic label to each cell in the local map representing the likelihood it belonged to each of four classes: grass, sidewalk, street, and unknown. Experimental results demonstrating autonomous sidewalk following for several hundred meters highlight the potential utility of the approach.

The remainder of this paper is organized as follows: after discussing related works, we introduce our perception platform in Section III. Next, we detail the map and model generation process in Section IV. In Section V we benchmark our model against a neural network implementation to assess its efficacy, and present the results of several autonomous experiments. We conclude with a discussion of results and future work in Section VI.

II. Related Work

Current approaches to terrain classification are diverse, and employ a wide array of hardware and algorithms to distinguish

terrain types. Typical methods can be divided into three groups: vibration-, vision-, and geometry-based.

Vibration-based approaches use an inertial measurement unit (IMU) to measure the accelerations a robot experiences during its operation. Typically, machine learning techniques compare inertial readings to data learned during a training period to classify terrain [9], [10]. The biggest shortcoming of this method is that the robot can only classify terrain it is currently on, thus rendering vibration-based classification unsuitable for avoiding undesirable terrain.

Vision-based techniques use a camera to capture information about the ground ahead of the robot. Color and texture information is analyzed and again compared against data learned in a training phase [11], [12]. However lighting and weather conditions can render the training set ineffectual.

Most recently, exteroceptive sensors such as 3D LIDARs have been used to classify terrain using range and geometry information [13], [14]. With these techniques, 3D point clouds are composed into a scene and geometric features are used to draw inferences about surface type. Terrain classification using 3D laser remission has been explored in [15] and [16]. Remission-based classifiers for active emitters such as LIDAR are exceptional as they are less dependent on ambient lighting conditions. Both Saitoh *et al.* and Wurm *et al.* use remission to discern street from grass. While Saitoh *et al.* report only 1% false positives using an online learning approach, their experimental data show significant false negative errors. The work of Wurm *et al.* is most similar to ours. They use range, bearing, and intensity as features to train a SVM, and demonstrate over 99% accuracy in both false positives and negatives. However, while they find that street and grass classes are easily separable by range, bearing and intensity (our work agrees with this), our findings indicate overlap when an additional sidewalk class is added. As a result, we propose a stochastic approach.

III. Developmental Platform

The SWS used in this work is shown at Figure 1a. It is based upon an Invacare M91 Pronto electric power wheelchair, but upgraded with the ATRS' Motion Control Module (MCM) to enable autonomous control [5]. For exteroceptive sensing, the platform integrates a pair of IFM O3D200 3D cameras highlighted in Figure 1b [17]. Unlike 3D sensors such as the Microsoft Kinect, the IFM can operate in direct sunlight. Moreover, the it provides both range and intensity measurements. The latter is indicative of the reflectivity of objects in the scene, which is crucial for our application. The IFM is compact and relatively low cost (less than $1,500 US). However, it has a relatively low resolution (50 $\times$ 64 pixels), a narrow field of view (30° $\times$ 40°), and limited effective range ($\approx$ 5 meters in our application). Despite these limitations, we could find no better 3D sensor available on the market at this price point. Note our SWS integrates a pair these sensors to augment the horizontal field of view.

Motor control for the SWS was accomplished with an onboard embedded computer. A software-based PID controlled

(a) Smart Wheelchair System (b) IFM O3D200 3D Camera

Fig. 1: The SWS platform (left) integrates two IFM O3D200 3D cameras for exteroceptive sensing, high-resolution encoders, GPS, and an inertial measurement unit. The IFM sensors (right) measure both range and reflectivity of objects in the environment.

the individual wheel velocities using feedback from high resolution quadrature encoders (4,096 CPR), and a Microstrain 3DM-GX1 inertial measurement unit provided gyro corrections for improved odometry performance. With the exception of motor control, all processing was done by a laptop computer with a 1.6 GHz Intel 720QM processor. The software architecture was based on Willow Garage's Robot Operating System (ROS) [18].

IV. Terrain Map Generation

The primary objective of our terrain classification approach is to build a map of urban terrain, which includes street, sidewalk, and grass terrain types. This additional semantic information has the potential to not only support terrain-based navigation, but also to enhance existing perception algorithms that rely upon geometry information alone. In this section, we discuss a discrete probabilistic classification model that associates a stochastic label with an observation, which then enables the construction of a terrain occupancy map.

The sensor employed in this work was an IFM Efector O3D200 3D camera system. It provides 3D range measurements, which enable the SWS to see far enough to make decisions based upon geometry; and remission measurements, which provide a metric to differentiate materials. The chosen feature space for terrain classification was range ρ, bearing angle ϕ, and remission value γ based on [15]. Thus the i^{th} observation in this space at time t is $\mathbf{z}_{i,t} = [\rho_i, \phi_i, \gamma_i]^\top$. Range and bearing angle were defined as

$$\rho_i = \sqrt{x_i^2 + y_i^2 + z_i^2} \tag{1}$$

$$\phi_i = \arctan 2(y_i, x_i) \tag{2}$$

where (x_i, y_i, z_i) is the i^{th} point in the camera image. Equation 2 assumes positive x is in front of the camera, positive y is to the left. Remission values were not normalized and read directly from raw scan data.

Figure 2a shows a sample of the relationship between these features for three different surface classes: grass in green, sidewalk in blue, and street in red. From [15], we expected grass and street to be easily separable in this feature space, and this was indeed the case. However, as illustrated in Figure 2a, there was overlap between sidewalk and grass classes. Hence, neither a straight-forward hyperplane classifier nor adaptive thresholding as applied in [15] and [16] were appropriate in this feature space for these three classes. Instead we used a probabilistic approach to apply a stochastic label to an observation in the feature space. We define a stochastic label Θ on an observation as

$$\Theta_{i,t} = \{p(\theta_1|\mathbf{z}_{i,t}), \ldots, p(\theta_M|\mathbf{z}_{i,t})\} \tag{3}$$

where $p(\theta_j|\mathbf{z}_{i,t})$ is the probability of observing class $j \in \{1, \ldots, J\}$ conditioned on the observation $\mathbf{z}_{i,t}$.

A. Terrain Map

Our representation for the terrain map was a 2D occupancy grid, where every grid cell held the probability for belonging to each terrain class. We decomposed the global map into a grid of cells M, where each map cell $m_k \in M$ contained a stochastic label that was updated as observations (transformed to the world frame) were made in that cell. For each new observation, this label was updated according to the update rule derived below.

We start with the standard Bayes Filter update rule under the Markov assumption that the current observation, conditioned on the observed class, is independent of any previous observations

$$p(\theta_{k,j}|\mathbf{z}_{1:t}) = \frac{p(\mathbf{z}_{i,t}|\theta_{k,j})p(\theta_{k,j}|\mathbf{z}_{1:t-1})}{p(\mathbf{z}_{i,t}|\mathbf{z}_{1:t-1})}. \tag{4}$$

By Bayes' Law we have

$$p(\mathbf{z}_{i,t}|\theta_{k,j}) = \frac{p(\theta_{k,j}|\mathbf{z}_{i,t})p(\mathbf{z}_{i,t})}{p(\theta_{k,j})}. \tag{5}$$

Substituting this into Eq. 4 leads to

$$p(\theta_{k,j}|\mathbf{z}_{1:t}) = \frac{p(\theta_{k,j}|\mathbf{z}_{i,t})p(\mathbf{z}_{i,t})p(\theta_{k,j}|\mathbf{z}_{1:t-1})}{p(\theta_{k,j})p(\mathbf{z}_{i,t}|\mathbf{z}_{1:t-1})}, \tag{6}$$

and similarly

$$p(\neg\theta_{k,j}|\mathbf{z}_{1:t}) = \frac{p(\neg\theta_{k,j}|\mathbf{z}_{i,t})p(\mathbf{z}_{i,t})p(\neg\theta_{k,j}|\mathbf{z}_{1:t-1})}{p(\neg\theta_{k,j})p(\mathbf{z}_{i,t}|\mathbf{z}_{1:t-1})}. \tag{7}$$

Taking the ratio of Eqs. 6 and 7 leads to

$$\frac{p(\theta_{k,j}|\mathbf{z}_{1:t})}{p(\neg\theta_{k,j}|\mathbf{z}_{1:t})} = \frac{p(\theta_{k,j}|\mathbf{z}_{i,t})p(\theta_{k,j}|\mathbf{z}_{1:t-1})p(\neg\theta_{k,j})}{p(\theta_{k,j})p(\neg\theta_{k,j}|\mathbf{z}_{i,t})p(\neg\theta_{k,j}|\mathbf{z}_{1:t-1})}, \tag{8}$$

which, by taking its logarithm, can be tranformed to the log-odds update formula

$$\begin{aligned}\log\frac{p(\theta_{k,j}|\mathbf{z}_{1:t})}{p(\neg\theta_{k,j}|\mathbf{z}_{1:t})} = &\log\frac{p(\theta_{k,j}|\mathbf{z}_{1:t-1})}{p(\neg\theta_{k,j}|\mathbf{z}_{1:t-1})} + \\ &\log\frac{p(\theta_{k,j}|\mathbf{z}_{i,t})}{p(\neg\theta_{k,j}|\mathbf{z}_{i,t})} - \log\frac{p(\theta_{k,j})}{p(\neg\theta_{k,j})}.\end{aligned} \tag{9}$$

This derivation is examined in greater detail in [**?**]. Thus to update a cell we only needed to calculate the cell prior $p(\theta_{k,j})$ and the inverse sensor model $p(\theta_{k,j}|\mathbf{z}_{i,t})$. This update was performed J times – once for each terrain class – for every observation in the map cell, and the resulting updated probabilities were finally normalized.

B. Discrete Stochastic Model

In the previous section, we showed that calculating the inverse sensor model is necessary to update the occupancy grid. In this section, we descibe the discrete probabilistic model we used to estimate the probability $p(\theta_{k,j}|\mathbf{z}_{i,t})$.

We constructed a probability mass function (PMF) to estimate the probability that a new sensor measurement belongs to each terrain class. The feature space was decomposed into $E \times F \times G$ bins, with edges according to

$$\begin{aligned}\rho_e &= \rho_{min} + \frac{e}{E}(\rho_{max} - \rho_{min}) \qquad e = 0, 1, \ldots, E \\ \phi_f &= \phi_{min} + \frac{f}{F}(\phi_{max} - \phi_{min}) \qquad f = 0, 1, \ldots, F \\ \gamma_g &= \gamma_{min} + \frac{g}{G}(\gamma_{max} - \gamma_{min}) \qquad g = 0, 1, \ldots, G\end{aligned}$$

Then $b_{e,f,g}$ was a bin in the feature space with edges ρ_e, ϕ_f, and γ_g. Let $T = \{\{\mathbf{z}_1, \theta_1\}, \ldots, \{\mathbf{z}_N, \theta_N\}\}$ be a set of N training samples, where each $\mathbf{z}_n$ is a real measurement in the feature space generated by the 3D LIDAR, and each $\theta_n \in \{1, 2, \ldots, J\}$ is an index for the true class to which $\mathbf{z}_n$ belongs. Elements of T were assigned to a bin if the Euclidean distance from the point to the midpoint of $b_{e,f,g}$ was minimal over all bins.

The probability $p(\theta_{k,j}|\mathbf{z}_i)$ was calculated as follows. Let the set of training samples belonging to class j be

$$x_j = \{\mathbf{z}_n|\theta_n = j\} \quad \forall n \in N \quad \forall j \in J \tag{10}$$

and $x_{j,e,f,g} \subset x_j$ is the set of all training points belonging to class j in bin $b_{e,f,g}$. Then the un-normalized probability that $\mathbf{z}_i$ in bin $b_{e,f,g}$ belongs to class j is then

$$p_j = \frac{|x_{j,e,f,g}|}{|x_j|} + p_{j_0} \quad \forall j \in J \tag{11}$$

where p_{j_0} is an adjustable parameter that characterizes unsampled members of class j in the cell. We normalize the cell probability to arrive at the final inverse sensor model

$$p(\theta_j|z_{i,t}) = \frac{p_j}{\sum_{j=1}^{J} p_j + \alpha} \quad \forall j \in J \tag{12}$$

$$p(\theta_{|J|+1}|z_{i,t}) = 1 - \sum_{j=1}^{J} p(\theta_j|z_t) \tag{13}$$

where α is an adjustable parameter that characterizes unknown classes in the cell.

An example of this stochastic model is depicted in Figure 2b. Probabilities are represented by color in the form of a (red,green,blue) triplet, where the color components are proportional to the respective class probabilities. This is especially evident at the border between sidewalk (blue) and grass (green).

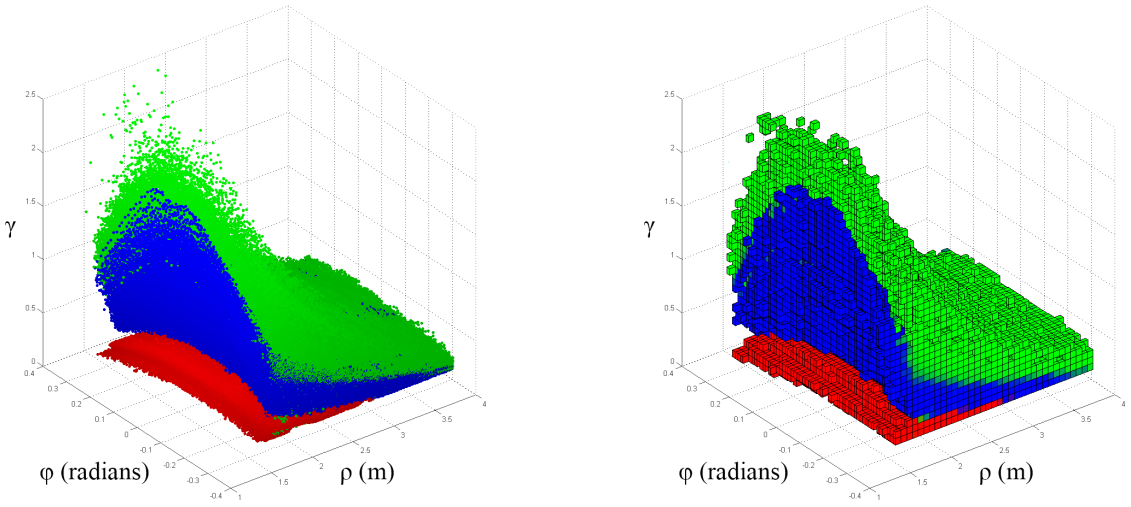

(a) Raw scans plotted in the feature space

(b) Discrete stochastic model

Fig. 2: In the left figure, remission data is plotted as a function of scan distance ρ and scan angle ϕ for three different material classes: grass in green (top), sidewalk in blue (middle) and street in red (bottom). The right figure is the stochastic model of this scan data, with bin color indicating class probability.

C. Generating the Sensor Model

Ground truth training data was collected using the IFM sensors mounted to the wheelchair platform described in Section III. The wheelchair was driven manually on several different representative surfaces for each of the three surface classes. Selected surfaces included old street, new street, old sidewalk, new sidewalk, overgrown grass, and freshly cut grass. To account for different lighting conditions, data was collected on each surface in bright sunlight, shade, and at night. To simplify the labeling process, the driven paths were chosen to ensure that all ground points belonged to the same class. Because sensor scans contain more than just ground points, we segmented and extracted the ground plane using RANSAC [19] with a planar model of the form

$$ax + by + cz = 0. \tag{14}$$

Outlier scan points were discarded, while inlier points were labeled with the associated surface class. In cases where the ground plane contained points from multiple classes, the scans were labeled manually by plotting the points in the feature space and visually estimating a boundary line between class regions.

In total 10,374,473 ground truth points were labeled, 5,038,779 of which were generated by the left IFM, and 5,335,694 of which were generated by the right IFM. Of these, 1,230,642 points (11.9%) were labeled as street; 6,082,385 points (58.6%) were labeled as sidewalk; and 3,061,446 points (29.5%) were labeled as grass. The number of points labeled for each class is indicative of their frequency of occurrence in the environment in which the data was collected; that is to say, that sidewalk is the most prevalent class, followed by grass, with street occurring relatively infrequently. These points were used to generate two models (as described in Section IV-B) respective to the IFM by which the points were generated. For both models model, the following parameters were used: $|E| = 100$, $|F| = 53$, $|G| = 100$, $p_{j0} = .05 \quad \forall j \in \{1, 2, 3\}$, and $\alpha = .1$. These values were arrived at experimentally.

V. Experimental Results

A. Benchmarking the Classifier

To assess the performance of the stochastic classifier, the entire ground-truth dataset from Section IV-C was evaluated. As an additional benchmark, a separate neural network classifier was also implemented for comparison purposes. Two neural networks were trained, one for each sensor, with data generated by that sensor in the following manner: A feedforward backpropagation neural network was used with a single hidden layer and 10 hidden neurons [20]. To train the network, ground truth data was partitioned into two sets: a training set and a testing set. From each class, 600,000 points were randomly selected for training using stratified holdout, for a total of 1,800,000 points in the neural network training set. To prevent overfitting the network to the training set, the training set was partitioned to create a validation set of data, which was evaluated continually during the training process. Network training terminated when the mean square error (MSE) of the validation set ceased decreasing. With the two networks generated, the entire set of ground truth data was classified by their respective models. The maximum likelihood was evaluated on the neural network classification results, and a confusion matrix was generated for each neural network.

The same subset of training points used to train the neural network was also used to build two stochastic models as detailed in Section IV-C, again one for each sensor. Again, the

Classifier	Blacktop	Sidewalk	Grass	All
Neural Network	99.8%	96.6%	94.5%	96.4%
Stochastic Model	100.0%	97.8%	95.2%	97.3%
Difference	0.2%	1.2%	0.7%	0.9%

TABLE I: Summary of benchmark classification performance.

Fig. 3: (Left) Satellite image of the path network used for the sidewalk following experiments. The route for a single trial is highlighted. (Right) Terrain classification results. The delineation between the sidewalk path (blue) and grass border (green) is clearly evident. Also the small black patch in the bottom right of the figure indicates an unknown class, which was a mulch bed. The distance traveled by the SWS was 210 meters.

entire ground-truth dataset was classified with these models, and maximum likelihood was used to apply a single label to each point. Aggregate classification results for the stochastic classifiers and the neural networks are shown at Table I as a confusion matrix. In total, the neural network correctly classified 96.5% of labeled data, while the stochastic approach correctly classified 97.3% of training data.

These results indicated slightly better performance from the stochastic classifier. More significantly, whereas the neural network took considerable time to train, the stochastic model can be generated in low-order polynomial time $O(N^d)$, where d is the dimension of the feature space and N is the size of the training set. In our implementation, training the neural network took over an hour on a quad core 3.0 GHz CPU with only a tenth of the total dataset. Training on the entire dataset of 10 million points was not attempted due to hardware limitations. In contrast, the stochastic model was generated in a fraction of a second, even on the entire 10 million point dataset. Further, updating a cell's weight in the stochastic model was a simple constant time operation. The net result is that the proposed stochastic approach can support real-time terrain classification, as well as online adaptation of the terrain models if required.

B. Terrain-based Sidewalk Following

Following benchmark testing, we evaluated the performance of the proposed terrain classifier in a sidewalk following task. In these experiments, the SWS was operated on sidewalk lined with grass. This test configuration was chosen since it represented a "worst case" pairing for our terrain set, as these two classes exhibited overlap in the chosen feature space as discussed in Section IV.

The SWS employed a two-dimensional occupancy grid as its local map M with a cell resolution of 5 cm. Measurements from the two IFM sensors were then processed separately. Points associated with the ground plane in each scan were then mapped to grid cells, and the respective probabilities of each cell being associated with a given terrain class were updated in accordance with Section IV. The prior probabilities for each class were uniform across all map cells. These probabilities were $p(\theta_1) = .9$, $p(\theta_2) = .6$, $p(\theta_3) = .8$, and $p(\theta_4) = .2$ for street, sidewalk, grass, and unknown respectively, and were arrived at experimentally.

For local planning, the SWS employed a sample-based planner on the input space of its linear and angular velocities (v, ω). The planning loop ran at 5 Hz, and the planning horizon was 2 seconds. The quality of the resulting trajectory T was evaluated using a cost function of the form $C(T, M) = C_{obst} + C_v + C_\omega$ where C_{obst} was related to its distance to obstacles, C_v was inversely proportional to v^2 to encourage forward progress, and C_ω was proportional to $|\omega|$ to discourage or encourage vehicle turning depending upon controller mode.

To effect the desired sidewalk following behavior, cells classified as grass were modeled as obstacles by the local planner. In initial experiments, this would often result in no feasible trajectories being found as a single misclassified sidewalk cell would map to an obstacle in M and block the path. This would occur at times when even a single leaf was present on the sidewalk. To mitigate this, a statistical outlier filter was used and dramatically improved the homogeneity of the classification result.

Three trials were subsequently conducted covering a distance of 675 meters. In each, the wheelchair was operated in "supervisor" mode, where navigation was completely autonomous unless interrupted by the user. In these experiments, user interrupts were limited to left and right turn inputs to the global planner so that the desired sidewalk route would be followed. The path and segmentation results from a representative trial are at Figure 3 (right). A significantly more informative video of a different trial can also be viewed at http://youtu.be/1RQLNKz1fQE. The net result was that using terrain information alone, the SWS was able to reliably follow the sidewalk with only higher level user intervention.

While we were initially satisfied with these results, subsequent testing identified environmental conditions that would

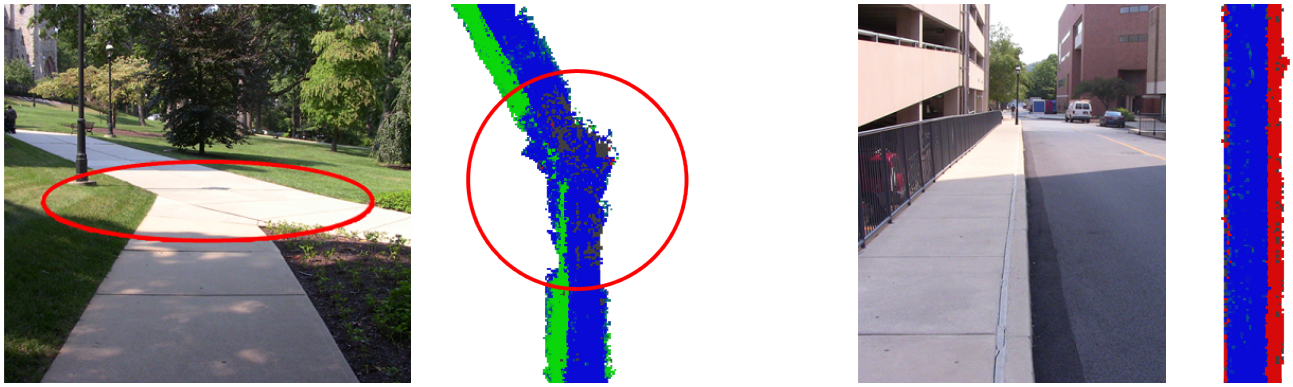

Fig. 4: (Left) Grass misclassified as sidewalk under bright sunlight conditions and with light colored sidewalk. No such errors were observed with the street class (shown in red), as illustrated on the right. Black cells in the terrain map indicate an unknown class.

exhibit consistent misclassification between the grass and sidewalk classes. This occurred under full sunlight with light colored sidewalk. An example is is illustrated at Figure 4 (left), with the misclassified regions circled in red. Note that the shaded grass regions are classified correctly. The cause of the misclassification appears to be associated with the IFM camera auto-clipping the specified exposure times due to high scene brightness, but this still needs to be verified. So while the IFM still produces range measurements, reflectivity measurements are corrupted. Note no such errors were observed with the sidewalk/street class pairings due to the distance between the two classes in the chosen feature space. This is illustrated in Figure 4 (right).

VI. Discussion

In this work, we investigated the potential for using intensity based terrain classification to augment the perception capabilities of a smart wheelchair system (SWS) operating in an urban environment. In benchmark testing, the proposed stochastic approach out-performed a neural network classifier without the extensive training requirements. The net result is a terrain classification approach suitable for real-time operations, as well as online adaptation of the terrain models as required.

We demonstrated the potential utility of terrain classification through a sidewalk following task. In a series of experiments, the SWS was able to discriminate between grass and sidewalk classes, and autonomously follow the sidewalk over distances of hundreds of meters solely from the class assignment. We acknowledge that these results are preliminary, and have already identified instances of class overlap between grass and sidewalk that would inhibit the sidewalk following task under certain conditions (*e.g.*, full sunlight on light colored sidewalk). Regardless, we should emphasize that we do not expect terrain classification alone to be sufficient for navigation. Instead, we see the strength of this additional semantic information in augmenting perception algorithms that currently rely solely upon geometry. For example, we have had success using iterative re-weighted least squares (IRLS) for ground plane segmentation as opposed to the more common RANSAC based approach [21]. In the least squares formulation, weights assigned to the individual measurements are a function of the distance to the ground plane recovered in the previous time step. An improved IRLS formulation might also employ the probabilities associated with terrain classification into the weighting, so that points identified as having a high probability of being sidewalk would be given greater consideration over those associated with the grass and street classes. We expect such an approach would provide a more accurate reconstruction of the ground plane associated with the sidewalk than otherwise possible using geometric information alone. Other uses that leverage both the range and intensity measurements provided by the 3D camera can be imagined, and a stochastic formulation provides a natural representation for merging this information.

Acknowledgments

This work was supported by the National Science Foundation (NSF) Robust Intelligence Program under Grant No. 0844585. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF.

References

[1] R. Simpson, "Smart wheelchairs: A literature review," *Journal of Rehabilitation Research and Development*, vol. 42, no. 4, pp. 423–436, Jul/Aug 2005.

[2] S. Hemachandra, T. Kollar, N. Roy, and S. Teller, "Following and interpreting narrated guided tours," in *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011.

[3] R. Bostelman and J. Albus, "Sensor experiments to facilitate robot use in assistive environemnts," in *Proc. of the Int. Conf. on Pervasive Technologies Related to Assistive Environments*, Athens, Greece, July 2008.

[4] J. Xu, G. Grindle, B. Salatin, D. Ding, and R. A. Cooper, "Manipulability evaluation of the personal mobility and manipulation appliance (permma)," in *Int. Symp. on Quality of Life Technology*, Las Vegas, United States, June 2010.

[5] C. Gao, I. Hoffman, T. Miller, T. Panzarella, and J. Spletzer, "Autonomous docking of a smart wheelchair for the automated transport and retrieval system (atrs)," *Journal of Field Robotics*, vol. 25, no. 4-5, pp. 203–222, 2008.

[6] ATRS: Automated Transport and Retrieval System. Freedom Science, LLC. [Online]. Available: http://www.youtube.com/watch?v=1fPgIPgzIfcl

[7] C. Montella, T. Perkins, J. Spletzer, and M. Sands, "To the bookstore! autonomous wheelchair navigation in an urban environment," in *Proc. of the Int. Conf. on Field and Service Robotics (FSR)*, Matsushima, Japan, July 2012.

[8] A. R. Meyers, J. J. Anderson, D. R. Miller, K. Schipp, and H. Hoenig, "Barriers, facilitators, and access for wheelchair users: sbstantive and methodologic lessons from a pilot study of environmental effects," *Social Science & Medicine*, vol. 55, no. 8, pp. 1435–1446, 2002.

[9] C. Weiss, H. Froehlich, and A. Zell, "Vibration-based terrain classification using support vector machines," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Bejing, China, October 2006.

[10] E. M. DuPont, C. A. Moore, E. G. Collins, and E. Coyle, "Frequency response method for terrain classification in autonomous ground vehicles," *Autonomous Robots*, vol. 24, no. 4, pp. 337–347, 2008.

[11] P. Vernaza, B. Taskar, and D. D. Lee, "Online, self-supervized terrain classification via discriminatively trained submodular markov random fields," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Pasadena, United States, May 2008.

[12] A. Howard and H. Seraji, "Vision-based terrain characterization and traversability assessment," *Journal of Roboic Systems*, vol. 18, no. 10, pp. 577–587, 2001.

[13] J. Lalonde, N. Vandapel, D. F. Huber, A. Kapria, and M. Hebert, "Natrual terrain classification using 3-d ladar data," *Journal of Field Robotics*, vol. 23, no. 10, pp. 839–861, 2006.

[14] B. Siemiatkowska, J. Szklarsk, and M. Gnatowski, "Mobile robot navigation with use of semantic map constructed from 3d laser range scans," in *Proc. of the Int. Conf. on Intelligent Information Systems (IIS)*, Siedlce, Poland, June 2010, pp. 47–60.

[15] K. M. Wurm, R. Kuemmerle, C. Stachniss, and W. Burgard, "Improving robot navigation in structured outdoor environments by identifying vegetation from laser data," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, St. Louis, United States, October 2009.

[16] T. Saitoh and Y. Kuroda, "Online road surface analysis using laser remssion value in urban environmnets," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, October 2010.

[17] O3d200 3d image sensor product specification. IFM Efector Inc. [Online]. Available: http://ifm.com/ifmus/web/dsfs!O3D200.html

[18] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009. [Online]. Available: http://www.ros.org

[19] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[20] M. F. Moller, "A scaled conjugate gradient algorithm for fast supervised learning," *Neural Networks*, vol. 6, no. 4, pp. 525–533, 1993.

[21] J. Bohren, J. Derenick, T. Foote, J. Keller, A. Kushleyev, D. Lee, B. Satterfield, J. Spletzer, A. Stewart, and P. Vernaza, "Little Ben: The Ben Franklin Racing Team's Entry in the 2007 DARPA Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 9, pp. 598–614, Sep 2008.