

Terrain Classification with Markov Random Fields on fused Camera and 3D Laser Range Data

Marcel Häselich* Marc Arends* Dagmar Lang* Dietrich Paulus*

*Active Vision Group, AGAS Robotics, University of Koblenz-Landau, 56070 Koblenz, Germany

{mhaeselich, marends, dagmarlang, paulus}@uni-koblenz.de

Abstract—In this paper we consider the problem of interpreting the data of a 3D laser range finder. The surrounding terrain is segmented into a 2D grid where each cell can be an obstacle or negotiable region. A Markov random field models the relationships between neighboring terrain cells and classifies the whole surrounding terrain. This allows us to add context sensitive information to the grid cells where sensor noise or uncertainties could lead to false classification. Camera images provide a perfect complement to the laser range data because they add color and texture features to the point cloud. Therefore camera images are fused with the 3D points and the features from both sensors are considered for classification. We present a novel approach for online terrain classification from fused camera and laser range data by applying a Markov random field.

In our experiments we achieved a recall ratio of about 90% for detecting streets and obstacles and prove that our approach is fast enough to be used on an autonomous mobile robot in real time.

Index Terms—Multi-sensor fusion, terrain negotiability

I. INTRODUCTION

Autonomous navigation in unstructured environments is a current and challenging task in robotics. The mobile system needs a detailed interpretation of the surrounding terrain to avoid obstacles and to regard the negotiability of the surface. A modern 3D *Laser Range Finder* (LRF) provides a rich and thorough picture of the environment in form of 3D distance measurements. Considering a path planning algorithm, the data of the 3D LRF cannot be directly used because of their vast amount. Therefore, as a first step, a reduction of the huge point cloud is necessary and an efficient data structure is essential. Our work provides a direct extension to the terrain analysis by Neuhaus et al. [14], where a data structure was introduced that we use in our work.

The laser range measurements alone allow no differentiation between different surfaces. Therefore we calibrated three cameras to the LRF and can then access the fused data in one coordinate system. This allows us to determine the color and texture of the 3D points in the field of view of each camera. In unstructured environments, the classification of the terrain can be a problem due to sensor noise, varying density of the data, egomotion or percussions on rough terrain. For that reason we apply a Markov random field to add context-sensitive information to the terrain classification, which models the relationships in our data structure.

Our goal is to determine the negotiability of the surrounding terrain with a Markov random field in real time based on the sensors described in the following section.

This paper is organized as follows. First of all we present the deployed sensors and their layout in Sec. II. After discussing related work in Sec. III, we will briefly describe the Markov random field basics in Sec. IV. The algorithm is described in detail in Sec. V followed by the experiments in Sec. VI. Finally, we present the conclusion in Sec. VII.

II. HARDWARE



Fig. 1. Deployed sensors: A 3D laser *Velodyne HDL-64E S2* (left) and two different commercially available cameras (*Logitech HD Pro Webcam C910* on the upper right and *Philips SPC1300NC* on the lower right) are used to perceive the environment.

As shown in Fig. 1 we use a *Velodyne HDL-64E S2* [8] and two different camera models. The head of the Velodyne consists of 64 lasers which permanently gather data of the environment as the head spins at a frequency from up to 15 Hz around the upright axis. In doing so, the sensor produces a rich dataset of 1.8 million distance measurements per second. The data of one full rotation are accumulated into one point cloud. A *Logitech HD Pro Webcam C910* [18] is installed to the front and two *Philips SPC1300NC* [15] cameras are fixed on the left and the right side of the construction. The sensors are either be mounted on top of a 500 kg robot (autonomous driving) or on a car (recording of sensor data).

III. RELATED WORK

There exist various approaches to classify the terrain surrounding an autonomous mobile robot platform. Especially image or laser based strategies are wide spread when terrain negotiability information is needed.

The image based strategies either use a single, stereo or combined setup of digital and infrared cameras. Konolige et al. [10] and Alberts et al. [1] both use stereo vision approaches to maneuver a vehicle through unstructured environments. Stereo vision allows them an extraction of traversable regions from the camera video streams. Furthermore, Vernaza et al. [23] present a camera based terrain classification approach. Their approach uses a Markov random field that classifies image data of a stereo system into obstacles or ground regions for an autonomous robot.

Negative obstacles (non-negotiable regions underneath the ground level) present a difficult challenge in non-urban environments. Thermal infrared images have the characteristic that negative obstacles remain warmer than the surrounding terrain in the night. Rankin et al. [17] therefore combine thermal signatures and stereo range data to determine the terrain negotiability.

Laser based approaches either work with a 2D, a self-build rotating 2D or a 3D LRF, like the *Velodyne HDL-64E S2*. Wurm et al. [26] use the laser remission value of a 2D LRF on a pan-tilt unit to classify the surface terrain based on the resulting 3D scans. In this way, they can detect grass-like vegetation and prefer paved routes with their robot.

Another approach for terrain classification is presented by Wolf et al. [25]. Their robot uses a 2D LRF oriented to the ground, records data while driving and produces 3D maps using Hidden Markov models. The authors are able to differentiate flat areas from grass, gravel or other obstacles.

Vandapel et al. [22] segment 3D distance measurements and classify the segments into three different classes for terrain surface, clutter or wires. Their approach worked with a detailed stationary 3D sensor as well as on a mobile platform with a rotating 2D scanning mount.

Ye and Borenstein [27] present an algorithm for terrain mapping with a 2D LRF. Their LRF is mounted at a fixed angle to the ground in front of their robot and creates an elevation map while driving.

A color stereo camera and a 2D LRF are used by Manduchi et al. [12] to detect obstacles. The authors present a color-based classification system and an algorithm that analyses the laser data to discriminate between grass and obstacles.

Schenk and Csatho [20] fuse aerial images with 3D point clouds to construct surfaces of urban scenes. The surfaces are represented in a 3D object space coordinate system by patches that store the shape and the boundary of the corresponding surface region.

A stereo pair of digital cameras, an infrared camera and two 2D LRFs mounted on scanning mounts are used by Wellington et al. [24]. The authors apply multiple Markov random fields that interact through a hidden semi-Markov model that enforces a prior on vertical structures. Their results showed that including the neighborhood structure significantly improved obstacle classification.

Beyond terrain classification another research topic covers the simultaneous localization and mapping (SLAM) problem, where LRFs are used to process the 2D and 3D distance measurements to build maps [5, 13] of the environment. In contrast to mapping, we want to process the sensor data as

fast as possible online on the robot so that a path planning algorithm can directly access these data and perform tasks autonomously in real time.

All presented approaches have in common, that unstructured environments are described as challenging by the authors. Sensor noise, manifold and complex vegetation, rough terrain and concussions while driving require a solid approach to realize an adequate classification of the terrain. Under the circumstances described, we want to achieve the best classification possible with our sensor data. A statistical approach can handle sensor noise and uncertainties quite well and for this reason we chose a Markov random field, which will be described in the following section.

IV. MARKOV RANDOM FIELD BASICS

Markov random fields can be assumed to be a 2D expansion of Markov chains. This implies that a random state ω_{ij} at position (i, j) has a neighborhood \mathcal{N}_{ij} including states ω_λ with

$$\omega_{ij} \neq \omega_\lambda \quad (1)$$

$$\omega_\lambda \in \mathcal{N}_{ij} \quad (2)$$

The neighboring states can be located in different distances forming different cliques of neighbors. In our work, we examine the 8 nearest neighbors surrounding a random state in a 2D data representation.

Assuming that each state takes a value results in a configuration ω of a random field. A random field can be called Markov random field if the probability for a configuration takes a positive value

$$P(\omega) > 0 \quad (3)$$

and if a random state depends only on its neighbor states

$$P(\omega_{ij} | \omega_{-\omega_{ij}}) = P(\omega_{ij} | \omega_\lambda), \text{ with } \omega_\lambda \in \mathcal{N}_{ij} \quad (4)$$

where $\omega_{-\omega_{ij}}$ stands for $\omega \setminus \omega_{ij}$ the set ω without ω_{ij} . For easier computation and modeling Markov random fields are often used as a Gibbs random field. Markov random fields and Gibbs random fields can be seen equal, which is described by the *Hammersley-Clifford Theorem*. In a Gibbs random field the distribution of the random states corresponds to the Gibbs distribution, which is defined as

$$P(\omega) = \frac{1}{Z} \exp\left(-\frac{1}{T}U(\omega)\right) \quad (5)$$

where T is a temperature parameter, $U(\cdot)$ a function which calculates the energy E of a configuration and Z is a normalizing constant with

$$Z = \sum_{\omega} \exp\left(-\frac{1}{T}U(\omega)\right) \quad (6)$$

that calculates a sum over all possible configurations. The modeling of $U(\cdot)$ depends on the problem, which needs to be solved. Maximizing $P(\omega)$ can be achieved by finding a configuration ω that minimizes $U(\omega)$. For this purpose several sampling techniques exist, as it is too time consuming to test every possible configuration.

An introduction of Markov random fields for classification tasks is presented by Theodoridis and Koutroumbas [21] and for a detailed description and several possibilities for modeling with Markov random fields are discussed by Li [11].

V. MARKOV RANDOM FIELD APPLICATION

After an introduction of the Markov random fields, the following section details their application to our sensor data and their use in our framework. The section starts with the data structure and pre-classification and continues with a description of the features. Our Markov Model of the terrain and the details of the camera and LRF fusion follow right after. The section ends with a description of the feature parameter estimation for the Markov random fields classes.

A. PRINCIPAL COMPONENT ANALYSIS

As a first step, the large point cloud delivered by our sensor is subdivided into smaller chunks, which can then be analyzed with the help of the *principal component analysis* (PCA) [14]. We use a 2D equidistant grid, centered around the origin of the sensor, to subdivide the point cloud. Our grid has a size of 40×40 m and each of the cells is scaled approximately 50×50 cm. The output of the PCA is a pre-classified grid, where each cell can either be *flat*, *obstacle* or *unknown*. A detailed description of the PCA in form of a hierarchical implementation is provided by Neuhaus et al. [14].

B. FEATURES

Features are essential for classification tasks. We define a set of features which describe different terrain classes we use in our approach. This set can be divided into two categories. The first type of features is acquired from of the data gained by the LRF mounted on our robot. The second type consists of features computed from the camera images of the terrain.

The first laser-based feature we use, is the *roughness* f_r of a terrain cell. The roughness can be obtained by calculating the *local height disturbance* [14]. The roughness feature offers valuable clues about how uneven a terrain cell is. This information is helpful for distinguishing between roads and grass or acres, for example.

The second laser-based feature, the *height difference* f_h , is capable of differentiating between obstacles and negotiable terrain. This feature simply displays the distance between the height of the lowest and the highest laser point of a terrain cell and was presented by Happold et al. [6]. The usage of this feature is based on our assumption that the main characteristic of an obstacle is its height.

Under the assumption that texture differs for diverse types of terrain, we apply features acquired from image data. We exert the *second angular moment* f_{sam} , *variance* f_v and *inverse difference moment* f_{idm} texture feature. These are 3 easy to compute features out of the 14 texture calculations proposed by Haralick et al. [7]. We use another fast calculating feature, expressing the homogeneity f_{fh} of a texture, proposed by Knauer et al. [9]. This is based on the assumption that rough terrain has an inhomogeneous texture in contrast to the texture

of even terrain. In addition to texture features, the color f_c of a terrain cell is also used, which is fast and easy to compute.

The Combination of all features yields the feature vector $\mathbf{f} = (f_r, f_h, f_{sam}, f_v, f_{idm}, f_{fh}, f_c)$ that integrates in the terrain classification process.

C. MARKOV MODEL OF THE TERRAIN

Regarding the fact that flat terrain cells can either be easy or hard to pass for a robot, the two classes *road* and *rough* are used for differentiation. This leads to a total of four possible terrain-classes $\{unknown, road, rough, obstacle\}$. A Markov model of the terrain should respect the context-sensitivity of neighboring cells and the acquired features. For this reason we chose a Markov random field model used for image segmentation tasks, described by Deng et al. [2]. This model is able to fulfill both requirements and can easily be adapted to the terrain classification problem by regarding each terrain cell as random state in the Markov random field. Here the energy E of a Gibbs random field is calculated piecewise by splitting it up into a component $E_{\mathcal{N}}$ that models the relationship of the neighboring cells and another component E_f which models the computed features of a cell.

The neighborhood component $E_{\mathcal{N}_{ij}}$ of a cell C_{ij} at position (i, j) in a grid is defined as

$$E_{\mathcal{N}_{ij}} = \sum_{\lambda \in \mathcal{N}_{ij}} (\beta \cdot \delta(c_{ij}, c_{\lambda})) \quad (7)$$

where c_{ij} is the terrain class assigned to the cell C_{ij} and c_{λ} is the terrain class assigned to a cell C_{λ} , which is part of the neighborhood of C_{ij} . The function $\delta(\cdot)$ returns a value of -1 for $c_{ij} \neq c_{\lambda}$ and $+1$ for $c_{ij} = c_{\lambda}$. β is used to weight a neighbors impact according to its distance to the cell. In our approach, we observe the 8 closest neighbors to a cell where the value of β is fixed and adjusted to our environment.

The corresponding feature component $E_{f_{ij}}$ of a cell C_{ij} is based on the assumption that the features of a cell are Gauss distributed. As an energy its computation is defined as

$$E_{f_{ij}} = \sum_k \left(\frac{(f_{ijk} - \mu_{ijk})^2}{2\sigma_{ijk}^2} + \log(\sqrt{2\pi}\sigma_{ijk}) \right) \quad (8)$$

where f_{ijk} is the k -th feature of C_{ij} and μ_{ijk} and σ_{ijk} are the mean respectively the standard deviation of the k -th feature of the class c_{ij} assigned to a cell C_{ij} .

Combining the two calculations, the complete energy E_{ij} of a cell C_{ij} can be calculated as follows

$$E_{ij} = E_{\mathcal{N}_{ij}} + \alpha \cdot E_{f_{ij}} \quad (9)$$

where α is a weighting constant used to control the influence of the different energy types.

For classification the sum of all computed energies E_{ij} needs to be minimized, which leads to a maximization of the a posteriori probability of the labeling of the terrain. These energies can be minimized by finding a label for each cell which fits best for the computed features and the labels of the neighbor cells. We apply the Gibbs Sampler described by Geman and Geman [3] to solve this task.




Fig. 2. Example scene of an unstructured environment with three different representations (best viewed in color). In each representation the images from the left, front and right camera are shown in the upper part. The upper image shows the 3D point cloud provided by the *Velodyne HDL-64E S2*. The color of a point is determined by its height. On the image in the middle the LRF data are fused with all three cameras. This way color, texture and 3D position are available in one coordinate system. The lower image shows the result of the MRF application: classified terrain cells with negotiability information. The classes *unknown*, *road*, *rough* and *obstacle* are respectively visualized in black, gray, brown and red.

D. CAMERA AND LIDAR FUSION

Our framework grants the possibility to fuse laser points and image pixels. This is achieved by a scenegraph representation of the used hardware, which provides transformation matrices between different sensors and corresponding intrinsic parameters. Thus we can build a lookup table for each laser point, which returns the corresponding pixel. Using this table finding the image region spanned up by the measured laser points is only a minor problem. Further information how to calibrate a camera to a LRF are available in the literature [19].

E. ESTIMATION OF FEATURE PARAMETERS

The application of the described Markov random field model depends on the knowledge of the means and standard deviations of each feature for each class. Thus, we estimate these parameters by annotating terrain cells with labels out of the used set of classes, by hand. From the annotated terrain it is possible to calculate the needed parameters of the features. The laser feature parameters can be applied in different environments in contrast to the image feature parameters, which need to be estimated separately for each scenario, because of the different appearances of the terrain classes. This has the advantage that the estimated parameters are based on real data acquired from the environment, which are carefully classified by a human observer.

VI. EXPERIMENTAL RESULTS

In our experiments we compared our terrain classification results with a ground truth. The ground truth is acquired by a human observer, who annotates the terrain cells per hand. By comparing the results with their respective ground truth we calculate the true positive rate and the false positive rate for each terrain class. We evaluated the classification for different scenarios, where the relative appearances of the terrain classes differ. An example scene is shown in Fig.2 where the 3D data, the fused data and the Markov random field result are visualized. Table I shows the result for a scenario, where the robot drives on a road through fields and small to medium sized vegetation and Fig.3 the corresponding ROC curves. In this table the usage of laser based features is denoted by L, selected Haralick features by H, the homogeneity feature by FH and color by C. Furthermore TPR is the true positive rate, or recall rate and FPR the corresponding false positive rate. Table II shows the results for a forest scenario, where the class *rough* appears only at single terrain cells and not at larger regions. In this case, the results show that the detection of the less appearing class does not work properly. The reason is the modeling of the terrain, in which we assume that terrain cells of one class show a tendency to group and not to exist as single cells. Furthermore we observed that the usage of image features does not improve the quality of the results in our approach. A reason for this may be that texture and color varies more than we assumed, which leads to high values for the standard deviations of these features. Thus a high deviation from the mean value of such a feature does not change the computed energy value sufficiently to have impact on the classification.

The evaluation of the runtime, shown in Table III, shows that the usage of image features multiplies the needed calculation time. In this table PCA expresses the usage of the PCA based algorithm without the application of our Markov random field and MRF stands for the application of our Markov random field. Using laser features exclusively leads to a fast runtime and good classification results. In all experiments the fixed values α and β were respectively set to 0.4 and 0.8.

Class/value	L	L + H	L + FH	L + C
Road/TPR	91.049 %	91.914 %	90.432 %	91.049 %
Road/FPR	1.523 %	1.146 %	1.158 %	1.372 %
Rough/TPR	74.618 %	73.828 %	75.158 %	75.293 %
Rough/FPR	1.190 %	1.103 %	1.306 %	1.422 %
Obstacle/TPR	92.826 %	91.785 %	92.881 %	92.552 %
Obstacle/FPR	3.565 %	4.426 %	3.221 %	3.118 %

TABLE I

Class/value	L	L + H	L + FH	L + C
Road/TPR	95.197 %	91.747 %	95.021 %	94.981 %
Road/FPR	4.336 %	5.571 %	4.312 %	4.499 %
Rough/TPR	5.680 %	0.406 %	2.840 %	3.854 %
Rough/FPR	0.292 %	0.101 %	0.259 %	0.281 %
Obstacle/TPR	96.234 %	95.365 %	96.076 %	96.339 %
Obstacle/FPR	6.900 %	6.686 %	7.276 %	6.793 %

TABLE II

Table I and Table II shows our results for a rural and a forest scenario, respectively. The true positive rate (TPR) and the false positive rate (FPR) are presented for the three classes *Road*, *Rough* and *Obstacle*. The columns are separated by the used features: laser based (L), Haralick features (H), the homogeneity feature (FH) and the color information (C).

Mode	Mean	Std. dev.	Max	Min
PCA	9.832 ms	1.757 ms	16.145 ms	5.871 ms
L (MRF)	62.667 ms	7.357 ms	82.240 ms	48.787 ms
L+H (MRF)	744.917 ms	131.507 ms	993.987 ms	421.465 ms
L+FH (MRF)	141.379 ms	11.233 ms	188.420 ms	121.733 ms
L+C (MRF)	144.439 ms	12.566 ms	187.013 ms	117.326 ms

TABLE III

Table III displays the runtime results for our approach. All values represent long term measurements and result from the direct application to real sensor data, without preprocessing or any form of data reduction. The terms L, H, FH and C are equivalent to Table I and II. The runtime of the PCA is compared to the Markov random field application with different features from single sensor fused sensor data.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we presented an approach to process the fused data of a modern 3D LRF and three cameras with a Markov random field. We thereby integrate context sensitivity and different features from both laser and image data. The application of the Markov random field results in a classified 2D grid with terrain negotiability information for each cell. We annotated terrain cells in different scenarios to learn the class parameters of the features. The resulting grid forms a perfect data structure for an autonomous system in unstructured environments. We achieved a recall ratio of about 90% for detecting streets and obstacles.

Limitations of our approach reveal at higher speed. The robot platform we use has a maximum speed of 14 km/h, but

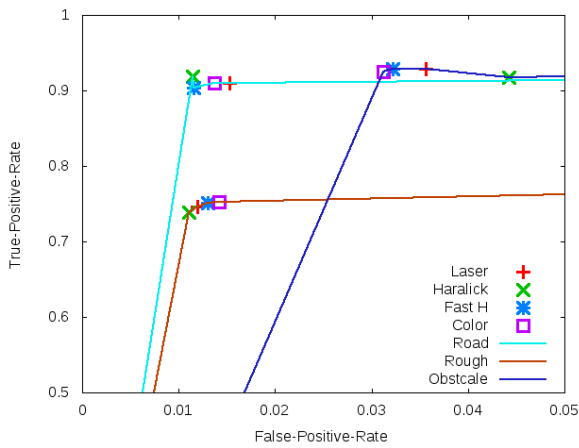


Fig. 3. ROC curves corresponding to the values of Table I. The three curves show the true positive rate in relation to the false positive rate for the three classes Road (cyan), Rough (brown) and Obstacle (blue)

the sensor data were also recorded while driving around with a car. For the usage of our approach on our autonomous system, the presented fusion is sufficient. In order to deal with very fast proper motion, the sensor calibration needs to concern the instant of time where the data was recorded. A system of precise time stamps should be able to solve this issue. Computation time of the image features is another difficult aspect. The image features need to be chosen wisely or a solution for speeding up this cost intensive part is needed.

In future work, we will work on an integration of time in the Markov random fields. By accessing previous classifications from the current, we hope to fill out spots with missing or sparse data. In example in regions with negative obstacles or in the immediate vicinity around our sensor. Another quite promising extension is a self-supervised learning of the feature parameters. In doing so, our approach can adapt to various situations and is able to improve during runtime if the features are relearned under certain circumstances. In addition, we will investigate the application of graph cuts [4] to the Markov random field data structure. The Markov random field and the fused data provide the opportunity to extract semantic information in real time from the sensor readings. This information can be used to create further information for 3D maps [16] and result in an semantic 3D map.

ACKNOWLEDGMENT

This work was partially funded by Wehrtechnische Dienststelle 51 (WTD), Koblenz, Germany.

REFERENCES

- [1] J. Alberts, D. Edwards, T. Soule, M. Anderson, and M. O'Rourke. Autonomous navigation of an unmanned ground vehicle in unstructured forest terrain. In *Proceedings of the ECSIS Symposium on Learning and Adaptive Behaviors for Robotic Systems*, pages 103–108, 2008.
- [2] H. Deng and D. A. Clausi. Unsupervised image segmentation using a simple mrf model with a new implementation scheme. *Pattern Recogn.*, 2004.
- [3] S. Geman and D. Geman. Stochastic relaxation, gibbs distribution, and bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1984.
- [4] D. M. Greig, B. T. Porteous, and A. H. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, 1989.
- [5] J. Guivant, E. Nebot, and S. Baiker. Autonomous navigation and map building using laser range sensors in outdoor applications. *Journal of Robotic Systems*, 2000.
- [6] M. Happold, M. Ollis, and N. Johnson. Enhancing supervised terrain classification with predictive unsupervised learning. In *Proceedings of Robotics: Science and Systems*, 2006.
- [7] R. M. Haralick, I. Dinstein, and K. Shanmugam. Textural features for image classification. In *Proceedings of IEEE Transactions on Systems, Man, and Cybernetics*, pages 610–621, 1973.
- [8] Velodyne Lidar Inc. Velodyne Lidar. <http://www.velodyne.com/lidar>, April 2011.
- [9] U. Knauer and B. Meffert. Fast computation of region homogeneity with application in a surveillance task. In *Proceedings of ISPRS Commission V Mid-Term Symposium Close Range Image Measurement Techniques*, 2010.
- [10] K. Konolige, M. Agrawal, R. C. Bolles, C. Cowan, M. Fischler, and B. Gerkey. Outdoor mapping and navigation using stereo vision. In *Proceedings of the International Symposium on Experimental Robotics*, pages 179–190, 2006.
- [11] S. Z. Li. *Markov Random Field Modeling in Computer Vision*. Springer, 2009.
- [12] R. Manduchi, A. Castano, A. Talukder, and L. Matthies. Obstacle detection and terrain classification for autonomous off-road navigation. *Autonomous Robots*, 2004.
- [13] M. Montemerlo and S. Thrun. Large-scale robotic 3-d mapping of urban structures. In *Experimental Robotics IX*, pages 141–150. Springer, 2006.
- [14] F. Neuhaus, D. Dillenberger, J. Pellenz, and D. Paulus. Terrain drivability analysis in 3d laser range data for autonomous robot navigation in unstructured environments. In *Proceedings of the IEEE International Conference on Emerging Technologies and Factory Automation*, pages 1686–1689, 2009.
- [15] Philips Electronics N.V. SPC1300NC. <http://www.p4c.philips.com/cgi-bin/dcbint/cpindex.pl?ctn=SPC1300NC/00>, April 2011.
- [16] J. Pellenz, D. Lang, F. Neuhaus, and D. Paulus. Real-time 3d mapping of rough terrain: A field report from disaster city. In *Proceedings of the IEEE International Workshop on Safety, Security and Rescue Robotics*, 2010.
- [17] A. Rankin, A. Huertas, and L. Matthies. Negative obstacle detection by thermal signature. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 906–913, 2003.
- [18] Logitech Europe S.A. HD Pro C910. <http://www.logitech.com/de-de/webcam-communications/webcams/devices/6816>, April 2011.
- [19] D. Scaramuzza, A. Harati, and R. Siegwart. Extrinsic self calibration of a camera and a 3d laser range finder from natural scenes. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4164–4169, 2007.
- [20] T. Schenk and B. Csatho. Fusing imagery and 3d point clouds for reconstructing visible surfaces of urban scenes. In *Proceedings of the IEEE GRSS/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas*, 2007.
- [21] S. Theodoridis and K. Koutroumbas. *Pattern Recognition*. Academic Press, 2009.
- [22] N. Vandapel, D. Huber, A. Kapuria, and M. Hebert. Natural terrain classification using 3-d lidar data. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 5117–5122, 2004.
- [23] P. Vernaza, B. Taskar, and D. D. Lee. Online, self-supervised terrain classification via discriminatively trained submodular markov random fields. In *IEEE International Conference on Robotics and Automation*, 2008.
- [24] C. Wellington, A. Courville, and A. Stentz. Interacting markov random fields for simultaneous terrain modeling and obstacle detection. In *Proceedings of Robotics Science and Systems*, 2005.
- [25] D. F. Wolf, G. Sukhatme, D. Fox, and W. Burgard. Autonomous terrain mapping and classification using hidden markov models. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2026–2031, 2005.
- [26] K. M. Wurm, R. Kümmerle, C. Stachniss, and W. Burgard. Improving robot navigation in structured outdoor environments by identifying vegetation from laser data. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1217–1222, 2009.
- [27] C. Ye and J. Borenstein. A new terrain mapping method for mobile robots obstacle negotiation. In *Proceedings of the UGV Technology Conference at the SPIE AeroSense Symposium*, pages 21–25, 2003.