Uncertainty-based Sensor Fusion of Range Data for Real-time Digital Elevation Mapping (RTDEM)

Lars B. Cremean†, Richard M. Murray
Division of Engineering and Applied Science
California Institute of Technology
†lars@caltech.edu

Submitted, 2005 IEEE International Conference on Robotics and Automation

***Abstract*— This paper introduces a new computationally inexpensive approach to perception and modeling of the environment that allows fusion of sensory range data of various types and fidelities while explicitly taking into account a complete description of uncertainty of the range measurements. This approach makes use of known sensor uncertainty models to create a single 2.5D digital elevation map whose accuracy is robust to sensor noise and spurious data. This approach is particularly suitable for real-time application in high speed and highly unstructured outdoor environments for which reasonably accurate and timely vehicle state estimates are available. Experimental results are presented in which LADAR range measurements and state estimates are combined according to this approach. We provide qualitative comparison to other classes of environment modeling.**

***Index Terms*— Terrain estimation, sensor fusion, 2.5D elevation mapping, range measurements**

I. Introduction

Autonomous navigation for mobile robots has much potential in civilian, commercial, military and space applications, but has not yet hit its stride in terms of demonstrating robust, real-time, high-speed operation in unstructured terrain for which there is no *a priori* terrain information available. Success in such an endeavor (and, indeed, in autonomous navigation in general), requires four fundamental contributions: *localization* within the environment, *perception and modeling* of the environment, *motion planning* through the environment and *execution* of planned motion.

The combined picture of these contributions have been treated in several contexts. Kelly and Stentz [1] have provided a thorough system-level overview of the requirements for navigation in rough terrain, as well as a dynamical-systems oriented approach for vehicle control in such a situation [2]. Lacroix et al. [3] have developed and demonstrated a comprehensive approach for navigating long distances in unknown environments, suitable for autonomous planetary exploration. Bellutta et al. [4] and Stentz et al. [5] demonstrated approaches to the terrain perception problem in particular, for the Demo III XUV and PerceptOR programs, respectively.

This paper is presented with all four of these fundamental components in mind (localization, perception, planning and execution), with a particular focus on the **perception and modeling** component, as applied in high-speed navigation in unstructured outdoor environments. Specifically, the goal addressed is the efficient and robust estimation of unstructured terrain given noisy state estimates and noisy range data. Since this terrain estimate is needed for real-time motion planning, issues of latency and throughput must be balanced with consideration for accuracy.

Fig. 1. Bob, Team Caltech's entry in the 2004 DARPA Grand Challenge, provides the test data for the results presented. Two stereovision camera pairs and a Sick LMS-2100 laser rangefinder are shown mounted above the cab; another Sick unit is mounted on the bumper. A Kalman Filter achieves state estimates from differential GPS and IMU input data.

Note that a natural division has and can be made within the modeling component between *terrain estimation* of the geometric properties of the environment and *terrain classification* of the surface and material properties of objects or regions in the environment. While the terrain classification problem has received various treatment in terms of both color and LADAR classification ([4], [6]), this paper is strictly concerned with the terrain estimation problem.

There have been several major approaches to the modeling (terrain estimation) component for autonomous navigation. These can be differentiated from each other by the manner in which the environment is chosen to be represented, in other words, the type of map that is used.

A large body of robotic literature represents the environment as a collection of distinct objects or landmarks in a *landmark map*, largely owing to the efficacy of

this representation in many solutions to the problem of Simultaneous Localization and Mapping (SLAM; for a thorough survey, especially as applied to indoor applications, see [7]). SLAM is particularly useful in applications where accurate localization estimates are not independently available.

A second popular type of environment representation for mobile robotics is the *occupancy grid*, which is a representation of the probability that each cell in the grid is occupied. Formal Bayesian Filtering methods have been well-developed for this type of map, and it is well-suited for navigation through structured terrain such as indoor environments. Extensions of the occupancy grid are closely related to the idea of evidence grids pioneered by Martin and Moravec [8], and these ideas have been extended to three dimensions in the form of *voxel maps*, which can provide much more accurate representations of environments at the cost of greater memory usage.

The type of map representation desired is highly dependent on the requirements of the specific application for which it is intended. For high-speed navigation, maintenance of complicated three-dimensional environment models is undesirable, as one of the co-requisite goals must be to maintain, process and evaluate the map with the minimum latency possible. On the other hand, navigation in unstructured terrain will limit the effectiveness of strictly 2D maps, since vital information about the surface geometry can be lost in these representations.

A suitable middle ground between these approaches, and the one that represents our framework for development, is the digital elevation map (DEM), where to each cell in a Cartesian grid is assigned the height estimate for that cell. This is a compact representation that is amenable to computationally fast implementation in terms of storage, access and evaluation. With this map representation, our problem statement reduces to the following: **Given a set of noisy range and vehicle state estimates, estimate the terrain surface elevation of the environment to provide an effective and efficient means of autonomous navigation**.

Several additional problem parameters will drive our approach. We will assume that reasonably accurate but noisy state estimates (3D location, pitch, roll and yaw) are available to us, and that we can co-register these state estimates with our range measurements, through the calibration parameters of our sensors. We further assume that our sensors are properly calibrated in terms of both the mounting parameters and intrinsic parameters of the sensor. The violation of this assumption would result in mis-registered maps; while there are methods to correct out calibration error, we leave the integration of these methods as future work. Finally, by virtue of adopting a digital elevation map approach, we are making the 2.5D world assumption; that is, we assume that for each (x, y) cell location there exists a unique elevation, and that the actual elevation can mathematically be described as a function, albeit not necessarily in closed-form. Violation of this assumption is tolerable, in such cases it may be impossible to show convergence, but formal proof of convergence is not a primary goal in this problem formulation.

Approaches to 2.5D terrain surface estimation in outdoor applications typically fall within one of two bipolar categories in terms of their treatment of uncertainty. The first approach neglects explicit account of uncertainty by averaging over the multiple measurements that fall within a given cell, and/or discarding outliers that fit poorly the other data associated with the corresponding terrain. The second approach typically constitutes a Bayesian formulation that constructs an expression for the probability of a 3D surface, conditional on the collection of range measurements [9], [10]. This approach, while most appealing because of its mathematical rigor, suffers a few drawbacks for our application. It's formulation requires an *a priori* model to which the measurements are compared, sometimes given as a parametric model whose parameters can be optimally estimated, neither of which are practical or desirable for navigation through unknown terrain. Finally, most methods that can accommodate this problem are not computationally efficient enough to process data as quickly and with as much throughput (on the order of thousands of measurements per second) as is required.

This paper, then, represents the "middle ground" between these two approaches. It provides a method to take range sensor uncertainty into explicit account in terrain estimation that is practical for real-time implementation at high speeds and data rates. Closely related approaches, in this respect, include those of Zhang [11], who estimates local parameters to fit a stereovision point cloud in a given region to a plane, Kelly and Stentz [2], who use the concept of a "scatter matrix" to represent the local geometric uncertainty in a grid, and Montemerlo and Thrun's recent approach [12] to accommodate sensor spatial resolution dependency on range.

The central contribution of this paper is the development of an approach for 2.5D terrain estimation which is amenable to sensor fusion, at the map level, of any number and variety of range sensors for which sensor models can be estimated. This approach makes explicit use of a sensor model and provides an efficient method of updating the grid cells for each range measurement. The main advantages of this approach are the generation of more complete, accurate and robust terrain estimates, along with a spatial measure of the *uncertainty* in the terrain estimate, which may guide new strategies for path planning or directing sensor attention. We exercise this new method on real field data of LADAR range measurements that are co-registered with vehicle state estimates obtained from a moving vehicle, and provide these raw data for further research in terrain estimation and autonomous navigation.

The paper is presented as follows. Section II provides a formulation of and motivation for our approach. Section III develops the mathematical preliminaries and derivation of each of the components of our approach. A description of the experiment run is given in section IV, and results

are provided in V. Section VI provides a summary of this work and description of current and future work.

II. Formulation of Approach

Traditional simultaneous localization and mapping solutions use both state and range measurements to update the map and location of the vehicle in the same update. State measurements (e.g. odometry) are used to update the locations of landmarks, and conversely, range measurements are associated with landmarks and used to update the estimate of the state of the vehicle. Our approach severs the latter connection, and propagates noisy but reliable state estimates through the usual geometric transformations to get a 3D description of the uncertainty of each measurement that encompasses the noise in the state as well as the range measurement. In this way, the data association problem is avoided altogether, since range measurements are not used to update the state of the vehicle.

The approach to estimate the terrain profile is outlined as follows:

1) Construct a sensor uncertainty model from estimated raw variances in each of the state estimate variables (easting, northing, altitude, pitch, roll, yaw) and from variances associated with the range measurement (range, azimuth, elevation).
2) Take a range measurement for which the transformation between the sensor frame and the vehicle frame is known. This range measurement is transformed through the uncertainty model in step 1 to achieve a 3D description of a probability density function (pdf) for the given measurement.
3) Choose a region of cells around the mean of the pdf calculated in step 2, and calculate updates for each of these cells.
4) Update the cells chosen in step 3 according to an appropriate set of update equations.

III. Mathematical Preliminaries

This section provides detailed derivation of one such implementation of the approach outlined in section II. Other implementations are possible and are likely to have their own advantages and disadvantages. Some remarks about the particular choice of implementation are provided in the subsections below. Subsection III-A describes an implementation of uncertainty model and range measurement pdf computation (steps 1 and 2), and subsections III-B and III-C describes a cell update method implementation (steps 3 and 4). These subsections represent the method used to achieve the results presented in section V.

A. Uncertainty Model

The following method is presented for computing the probability density function for a given range measurement. We refer to the pose of the vehicle at any given time as the collection of {easting, northing, altitude, pitch, roll, yaw}, denoted $\{x, y, z, \theta, \psi, \phi\}$, defined with respect to some inertial reference frame. We denote the variance in the estimate of each of these quantities as $\{\epsilon_x, \epsilon_y, \epsilon_z, \epsilon_\theta, \epsilon_\psi, \epsilon_\phi\}$.

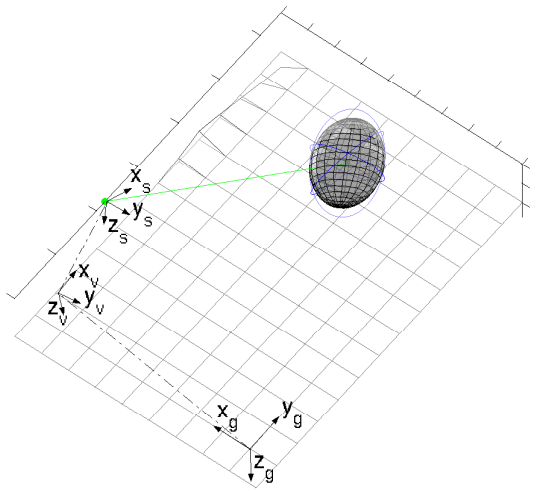

Fig. 2. The coordinate system conventions used in this paper include descriptions in the global, vehicle and sensor frame. Visualization of a measurement probability density function as a fixed-probability error ellipse. Scales are exaggerated for clarity.

These raw state variances can be provided, for example, from the internal state of a Kalman Filter state estimator, or they can be estimated offline. Let the measurement in the sensor frame be defined by its range, azimuth and elevation, denoted $\{\rho, \alpha, \beta\}$ and let the variances in these quantities be denoted by $\{\epsilon_\rho, \epsilon_\alpha, \epsilon_\beta\}$. We use in each of our coordinate frames the x-axis forward, y-axis right and z-axis down convention, and define $\theta = \psi = \phi = 0$ when the vehicle is flat and pointed north. These coordinate systems are presented for clarity in Fig. 2.

Our goal in this section is to derive a closed-form expression for the three-dimensional probability density function of a measurement, assuming each of the variance parameters defined above is given by Gaussian white noise. The location of a range measurement in sensor coordinates is given by

$$M_s = \begin{bmatrix} \rho \cos\alpha \cos\beta \\ \rho \cos\beta \sin\alpha \\ \rho \sin\beta \end{bmatrix}. \tag{1}$$

Applying the assumption of Gaussian noise (substituting $\rho \to \rho_0 + \epsilon_\rho$, $\alpha \to \alpha_0 + \epsilon_\alpha$, and $\beta \to \beta_0 + \epsilon_\beta$, and making a small angle approximation on the angle variances, yields a description of the noisy measurement in the sensor frame. After some algebra, this noisy measurement can be written as

$$M_s = M_{s0} + \epsilon M_s = \begin{bmatrix} \rho_0 c\alpha_0 c\beta_0 \\ \rho_0 s\alpha_0 c\beta_0 \\ \rho_0 s\beta_0 \end{bmatrix} + \ldots \tag{2}$$
$$\begin{bmatrix} c\alpha_0 c\beta_0 & -\rho_0 s\alpha_0 c\beta_0 & -\rho_0 c\alpha_0 s\beta_0 \\ \rho_0 c\alpha_0 c\beta_0 & s\alpha_0 c\beta_0 & -\rho_0 s\alpha_0 s\beta_0 \\ s\beta_0 & 0 & \rho_0 c\beta_0 \end{bmatrix} \begin{bmatrix} \epsilon_\rho \\ \epsilon_\alpha \\ \epsilon_\beta \end{bmatrix}$$

where the functions sin and cos are abbreviated. Note that the expression above is the combination of a nominal term equal to the measurement value and an uncertainty term

due to the variances in the sensor measurement. Based on (2), we can describe the probability density function in the sensor frame as centered around the nominal (measured) value with covariance given by $\Sigma = E(\epsilon M_s \epsilon M_s^T)$. Our goal, however, is to get the description of our measurement in the inertial frame in order to register it to our 2.5D map. This is done by transforming our measurement from the sensor frame to the vehicle frame, and then into the inertial frame. Noise in both of these transformations is accounted for by replacing the transformation variables with a nominal value plus additive noise, e.g. $x \rightarrow x_0 + \epsilon_x$ or $\phi \rightarrow \phi_0 + \epsilon_\phi$. Through this substitution, small angle linearization, and discarding of higher order terms, first order statistical expressions for the measurement in the vehicle and inertial frame can be obtained. Using the subscript syntax $_{vs}$ to indicate the transformation from the sensor to the vehicle, $_0$ to indicate the nominal measurement, and the prefix ϵ to indicate the first order statistic, the measurement in the vehicle frame can be written

$$M_v = R_{vs0} M_{s0} + R_{vs0} \epsilon M_s + \epsilon R_{vs} M_{s0} + T_{s0} + \epsilon T_s, \quad (3)$$

where R and T represent the rotation and translation of the associated transformation. Following the same procedure for transformation of the measurement into the inertial (global) frame yields a similar expression for the measurement in terms of (3),

$$M_g = R_{gv0} M_v + \epsilon R_{gv} (R_{vs0} M_{s0} + T_{s0}) + T_{v0} + \epsilon T_v. \quad (4)$$

This measurement description in the global frame is the sum of the nominal sensor measurement (transformed to the global frame) plus a first-order term that depends upon the combined variances from the sensor measurement itself as well as the variances which appear in both of the coordinate transformations. Its elaboration here is prohibited by space constraints, but it does constitute a closed-form expression consisting of additions and multiplications that lends to very fast computation (on the order of microseconds for a single measurement).

Through these noise propagation equations, each measurement can be described as a first-order probability density function with mean and covariance as derived from (4). The ellipsoid in Fig. 2 represents a fixed-probability error surface as computed from one such measurement description in the global frame.

B. Measurement Discretization

The explicit, computable expression for the measurement uncertainty resulting from the analysis of Section III-A will, in general, extend over multiple cells in a gridded map for a single measurement. It is therefore necessary to devise an appropriate way to determine which cells should be updated, and how they should be updated, for each measurement.

Subsection III-C describes the method we use to update a single cell given normally distributed cell input measurements z_k that are described by their mean z_m and variance σ_z. This section describes the way in which a single 3-D measurement and associated uncertainty model is converted into a number of cell input measurements.

Simplifying assumptions were made about the equations presented in Section III-A so that the uncertainty of a single measurement can be represented by first order statistics as a multivariate Gaussian. The shape and orientation of the Gaussian for each measurement depends on the direction of the sensor measurement as well as the variance in sensor range measurement and variances in vehicle position and orientation. The measurement is parameterized by a center $\mu \in \mathbb{R}^3$ and a covariance $\Sigma \in \mathbb{R}^{3\times3}$, and is represented by the equation

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{3/2}\sqrt{\det \Sigma}} \exp\left[-\frac{1}{2}(\mathbf{x}-\mu)^T \Sigma^{-1} (\mathbf{x}-\mu)\right] \quad (5)$$

where $p(\mathbf{x}) = p(x, y, z)$ is the probability that the measurement actually came from a surface at $\mathbf{x}$.

With an elevation grid representation of the environment, each cell C_{ij} in the grid can be reasonably assigned the height probability function

$$p_{ij}(z) = \iint_{C_{ij}} p(x, y, z)\ dy\ dx. \quad (6)$$

This function cannot be considered a probability density function, however, because the total integral of $p_{ij}(z)$ is not equal to 1. A normalization of this function is necessary to use it as a pdf for input to the Kalman Filter update equations of Section III-C. One option for normalization would be a scaling of the function by α_{ij} so that

$$\alpha_{ij} \int_{-\infty}^{\infty} p_{ij}(z)\ dz = 1. \quad (7)$$

The implementation of this method, however, has the significant drawback that the cells far from the center of a measurement will have variances that are similar to those of the cells near the center of the measurement. For cells far away from the measurement, the fact that $p_{ij}(z)$ is much smaller should correspond to a high variance associated with the measurement. Our solution, then, is to find the mean μ_{ij} of $p_{ij}(z)$ and then to normalize $p_{ij}(z)$ by setting the standard deviation of the normalized Gaussian pdf to

$$\sigma_{ij} = \frac{1}{p_{ij}(\mu)\sqrt{2\pi}}. \quad (8)$$

Cells close to the center of the measurement, therefore, will have higher value of $p_{ij}(\mu)$ and a lower variance. Cells far from the center of the measurement will have lower $p_{ij}(\mu)$ and hence a higher variance.

In a practical implementation, it is necessary to decide which cells are to be updated for any given measurement. There are options for how to specify this. One method for doing so for a Gaussian $p(\mathbf{x})$ is to update any cell whose center lies within a χ-confidence ellipse. Another is to update those cells for which the peak probability is greater than some threshold. A third is to say that we will update the cells whose centers lie within a specified geometric shape (rectangle, circle, ellipse). In the implementation

described in Section V, this last method was used for the sake of ease of implementation. Also for practical considerations, the integral in equation 6 was approximated by

$$p_{ij}(z) \approx p(x_i, y_j, z)\Delta_x\Delta_y \tag{9}$$

where the center of cell C_{ij} is (x_i, y_j) and the resolution of the grid is specified by Δ_x and Δ_y.

C. Cell Update Equations

The update equation that governs each cell is a Kalman Filter whose state is equal to the scalar height estimate for the cell. Since there are no independent dynamics associated with the height of a cell, the state propagation equations are simply identity, and the height estimate is purely a result of updates from measurements whose (x, y) coordinates land near that cell. In this manner, cells in the map receive a flurry of updates when a series of sensor measurements pass over it, but are otherwise unchanged. This results in an efficient means of updating cells during high-speed navigation.

An abbreviated version of the Kalman Filter measurement update equations (see for example [13]) is

$$K_k = P_{k-1}H^T(HP_{k-1}H^T + R)^{-1} \tag{10}$$

$$\hat{x}_k = \hat{x}_{k-1} + K_k(z_k - H\hat{x}_{k-1}) \tag{11}$$

$$P_k = (I - K_kH)P_{k-1} \tag{12}$$

In our case, the state vector is actually the scalar height of a cell, which is also the massaged form of our measurement, so the H term is equal to unity. The error covariance P is equal to the variance in the height estimate σ_h^2 and R is equal to the variance in the measurement height, σ_m^2. It can be shown from these equations that the update equations reduce to

$$\hat{x}_k = \frac{\sigma_h^2 z_m + \sigma_z^2 x}{\sigma_h^2 + \sigma_z^2} \tag{13}$$

$$P_k = \frac{\sigma_h^2\sigma_z^2}{\sigma_h^2 + \sigma_z^2} \tag{14}$$

These are the equations used to update the height and variance in height for the cells that are chosen to be updated according to the methods of section III-B.

IV. Description of Experiment

Experiments were run using Bob (Fig. 1), a instrumented and actuated sport-utility vehicle that served as Team Caltech's entry in the 2004 DARPA Grand Challenge. Bob is equipped with two pair of stereovision cameras and two 2D scanning laser rangefinders mounted horizontally. For tests of these new algorithms, simultaneous state and range data (from the cab-mounted LADAR unit only) was taken at approximately 5Hz and a manually controlled path was driven through a handcrafted course.

The course consisted of a flat dry lakebed with hand-placed obstacles at measured locations on the course. Included in the set of obstacles are three approximately 1m tall by 1m radius buckets, three coolers of approximate

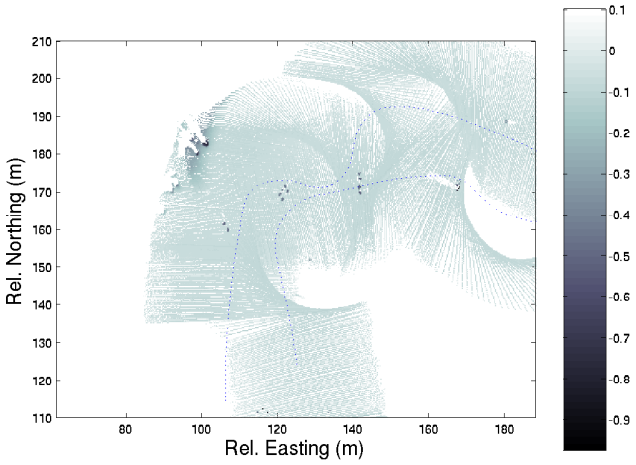

Fig. 3. The map resulting from the naive approach of replacing cell data with the height of the new measurement if the new measurement is higher than the old.

dimension 1m by 0.5m by 0.5m, and a calibration target of height and width 1m by 1m. Although the algorithms presented here are intended for high-speed operations, the experiments were run at moderate speeds in order to use the data to demonstrate the effectiveness of the algorithm under controlled conditions.

While demonstrated at moderate speeds, the algorithms developed based on sections II and III are able to run at very fast rates (processing the LADAR range measurements – 201 per scan at $5Hz$ scan rate – was demonstrated in real-time with a latency of a few tens of milliseconds on a 1.5GHz Pentium 4 processor). A thorough quantitative analysis of the algorithms performance at high speeds has not yet been done.

V. Results

The dotted line in Figs. 3 and 4 represents the path that was taken by Bob during one data run. The maps created in each of these figures are both 0.25m $\times$ 0.25m resolution, and represent the application of two different map estimation approaches. The first, naive method processed the data according to a method that neglected considerations of uncertainty in the measurement, and simply replaced the data in the map at the appropriate cell with the maximum of the measurement height and the current height in the cell, if any. The resultant map using this method is shown in Fig. 3. The second approach (Fig. 4) represents the result of the map update method presented in section II, with the cell containing the measurement mean and the eight cells immediately surrounding it receiving updates for each measurement.

The resulting terrain maps in Figs. 3 and 4 are coded with intensity proportional to height. Cells that are coded white represent no data assignments. Several qualitative comparisons can be made from these results. First, most strikingly, the maps provided by the new approach provided in this paper result in much more complete terrain maps

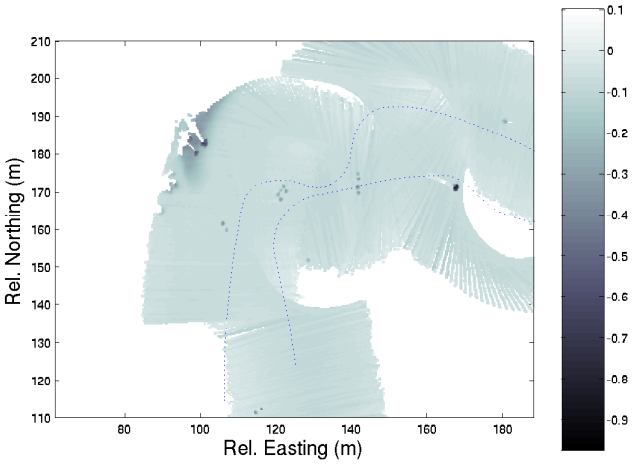

Fig. 4. The map resulting from the new approach, which updates a small region of cells for each measurement according to our update method.

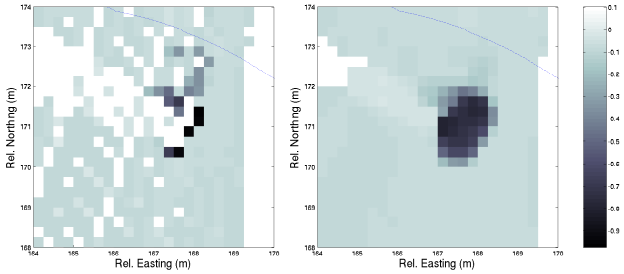

Fig. 5. A zoomed portion of Figs. 3 and 4 (on the left and right, respectively) near relative location $(170, 170)m$. Note that the sensor fusion method fills in the obstacle more fully and reduces the effect of spurious measurements.

than the simpler approach, meaning that many fewer cells remain assigned with no data. This is a direct result of updating a local region of cells for a given measurement. Secondly, this new approach tends to pick out obstacles more clearly than the naive approach, as seen for example in the zoomed image in Fig. 5. Third, the fused data method of creating the terrain maps has a smoothing effect on the terrain estimate. This is an expected effect of the algorithm.

Also note that for both methods of terrain estimation, registration errors are apparent for those obstacles that are passed more than once in the course of the experiment (those in the figures with less than about $150m$ relative easting). This effect is one disadvantage of not attempting to correct the state estimate with the range measurements, as is done in SLAM approaches. Application-specific considerations must be made when deciding whether this trade-off is desired.

VI. Summary and Future Work

The main contribution of this paper is to provide an approach to real-time 2.5D terrain estimation that explicitly uses sensor models in its formulation, including the effect of noisy but reliable state estimates. One such implementation of this approach was presented along with a qualitative analysis of its performance. It is expected that other implementations of this approach will be able to show marked improvement in terrain estimation for high-speed navigation through unknown and unstructured terrains.

This paper is intended as a baseline for research into computationally inexpensive methods of real-time digital elevation mapping approriate for high-speed navigation. It serves as a springboard to more mathematically formal approaches that are still amenable to processing at high speeds and high data rates. Future work includes development of such approaches, and providing rigorous quantitative analysis and comparison of different techniques.

Specific potential areas for future work include the development of data-driven sensor uncertainty models, as they may show marked improvement over the Gaussian assumption on the measurement model. We also intend to extend this general method to sensor fusion at the map level of multiple types of sensors, including combined stereovision and LADAR.

Acknowledgment

The authors would like to thank Dima Kogan, Kristo Kriechbaum, Shaunak Sen, and Abhishek Tiwari for course-based collaboration, and Joel Burdick for helpful discussions and advice.

References

[1] A. Kelly and A. Stentz, "An analysis of requirements for rough terrain autonomous mobility," *Autonomous Robots*, vol. 4, no. 4, December 1997.

[2] Alonzo Kelly and Anthony Stentz, "Rough terrain autonomous mobility – part 2: An active vision, predictive control approach," *Autonomous Robots*, , no. 5, pp. 163 – 198, May 1998.

[3] S. Lacroix, A. Mallet, D. Bonnafous, G. Bauzil, S. Fleury, M. Herrb, and R. Chatila, "Autonomous rover navigation on unknown terrains: Functions and integration," *International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 917–942, Oct-Nov. 2002.

[4] P. Bellutta, R. Manduchi, L. Matthies, K. Owens, and A. Rankin, "Terrain perception for demo III," in *Proc. of the Intelligent Vehicles Conference*, 2000.

[5] A. Stentz, A. Kelly, P. Rander, H. Herman, O. Amidi, R. Mandelbaum, G. Salgian, and J. Pedersen, "Real-time, multi-perspective perception for unmanned ground vehicles," in *Proc. of AUVSI*, 2003.

[6] R. Manduchi, A. Castano, A. Talukder, and L. Matthies, "Obstacle detection and terrain classification for autonomous off-road navigation," 2004, Accepted for publication.

[7] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel, Eds. Morgan Kaufmann, 2002, To appear.

[8] Martin C. Martin and Hans Moravec, "Robot evidence grids," Tech. Rep. CMU-RI-TR-96-06, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, March 1996.

[9] Ross T. Whitaker and Jens Gregor, "A maximum-likelihood surface estimator for dense range data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 10, October 2002.

[10] Ross T. Whitaker and Ernesto Lautaro Juarez-Valdes, "On the reconstruction of height functions and terrain maps from dense range data," *IEEE Transactions on Image Processing*, vol. 11, no. 7, July 2002.

[11] Zhengyou Zhang, "A stereovision system for a planetary rover: calibration, correlation, registration, and fusion," *Machine Vision and Applications*, vol. 10, no. 1, pp. 27–34, 1997.

[12] M. Montemerlo and S. Thrun, "A multi-resolution pyramid for outdoor robot terrain perception," in *Proceedings of the AAAI National Conference on Artificial Intelligence*, San Jose, CA, 2004, AAAI.

[13] Greg Welch and Gary Bishop, "An introduction to the kalman filter," Tech. Rep. TR 95-041, University of North Carolina, 1995.