Estimating the Heading of a Husky Mobile Robot with a LiDAR Compass based on Direction Maps*

Marc J. Gallant, Joshua A. Marshall, and Brian K. Lynch

***Abstract*— Without an absolute position sensor (e.g., GPS), an accurate heading estimate is necessary for proper localization of an autonomous unmanned vehicle or robot. This paper introduces direction maps (DMs), which represent the directions of only dominant surfaces of the vehicle's environment and can be created with negligible effort. Given an environment with reoccurring surface directions (e.g., walls, buildings, parked cars), lines extracted from laser scans can be matched with a DM to provide an extremely lightweight heading estimate that is shown, through experimentation, to drastically reduce the growth of heading errors. The algorithm was tested using a Husky A200 mobile robot in a warehouse environment over traverses hundreds of metres in length. When a simple *a priori* DM was provided, the resulting heading estimation showed virtually no error growth.**

***Keywords*— LiDAR compass, heading estimation, direction maps, localization**

I. INTRODUCTION

Mobile robot localization is the process of determining the position and orientation of a mobile robot relative to its environment. It is important in many common applications, such as planning, mapping, and control. Regardless of the methods used for localization, accurate determination of the heading of the robot is particularly important when interoceptive sensors are used. Motion models transform the outputs of these sensors into estimated changes of the robot pose. In this common scenario, the positional components of the motion model are usually tightly coupled with the heading estimate, meaning inaccuracies in heading estimation are quickly propagated into position errors.

As a result, relying on dead reckoning for heading estimation (e.g., using wheel encoders, gyroscopes) will eventually lead to substantial position errors. These errors occur regardless of the accuracy of the translation estimates provided by the interoceptive sensors. Absolute heading sensors (e.g., compass, sun sensor) can bound heading estimation errors, but are limited by their restrictive operating requirements (e.g., outdoors, no magnetic interference, etc.). Much like a compass uses the known direction of the Earth's magnetic field, or a sun sensor uses the known trajectory of the sun, an external reference with known orientation is needed to provide an absolute reference to the robot from which its heading can be absolutely determined. Often, this reference comes in the form of a map of the local environment.

*This work was supported through the Natural Science and Engineering Research council of Canada under CREATE grant 397985-2011 and by the NSERC Canadian Field Robotics Network (NCFRN).

The authors are with the Mining Systems Laboratory at Queen's University in Kingston, ON, Canada {m.gallant, joshua.marshall, lynch.brian}@queensu.ca

Providing or generating useful maps for robot localization is a non-trivial endeavour. If a map is provided beforehand, the data generated by exteroceptive sensors must be compatible with the map representation. Put differently, correspondences between sensor data and the map must be detectable (e.g., line extraction from laser data and line-based maps). To generalize, simultaneous localization and mapping (SLAM) can be used to create maps using the same types of sensors used for localization. However, SLAM is a complex problem involving computation requirements that grow as the environment is explored. If a separate mapping stage can be performed (i.e., sensor data of the environment is collected and processed offline), accurate maps can be generated by using state-of-the-art graph-based SLAM techniques [1]. These methods can even be used for real-time localization at the cost of substantial computation requirements [2].

This paper describes an extremely lightweight absolute heading estimation technique whereby *a priori* maps can be easily provided to the robot. These *direction maps* (DMs) are not intended to represent a geometric layout of the environment, but rather a compact representation of the directions of dominant local surfaces. These maps can be generated for any environment where the directions of observable surfaces occur in non-random sets (e.g., the walls of an indoor space, the exteriors of buildings and parked cars outdoors). Much like a compass observes the direction of the Earth's magnetic field, these surfaces are observed by a 2D scanning LiDAR, which uses their directions as an absolute reference from which the heading of the robot can be estimated. In this way, the sensor acts as a *LiDAR compass*. By combining this technique with an interoceptive sensor capable of accurately measuring translation (e.g., wheel encoders), fairly accurate localization can be achieved with minimal computation and without the need for a separate mapping stage.

A. Related Literature

Although DMs represent only the directions of lines in the environment (without positional information), it is important to note that the benefits of using a line-based map representation for localization and SLAM have been explored before. Cox [3] introduced one of the first implementations, which provided online localization given an *a priori* line-based map. Vandorpe *et al.* [4] introduced map building using geometric primitives (including line segments). More recent efforts have used line-based maps in SLAM implementations. Garulli *et al.* [5] demonstrated that localizing in SLAM-generated line-based maps can be

accurate. Nguyen *et al.* created the popular *OrthoSLAM* [6], which is an indoor line-based SLAM that takes advantage of an orthogonality assumption of the surfaces in the environment (e.g., indoor spaces with perpendicular walls). Kuo *et al.* [7] used the compactness of line-based maps and the orthogonality assumption to create a lightweight Rao-Blackwellized particle filter SLAM. Despite the computation reduction of line-based map representations, the requirement of positional information in the above implementations often require complex calculations to extract the uncertainty of the line parameters. Also, no concessions are made for areas in the environment that violate the orthogonality assumption.

Recently, a different approach of heading estimation using cameras has been explored, most commonly referred to as a visual compass (VC). A common approach is to unwrap sequential omnidirectional images and observe how simple extracted features are displaced as the robot rotates [8], [9]. An overview by Labrosse [8] of this method concluded that its dead-reckoning heading estimation performed similarly or better than inertial sensors in appropriate environments. Other forms of VCs do not necessarily require omnidirectional cameras and instead track the motion of specific vanishing points at far distances [10], [11], which work well in large, open environments. Finally, a recent VC implementation by Mariottini *el al.* [12] addresses the restrictions of many other VCs (e.g., computation, prior environmental knowledge, or calibration requirements) with an algorithm that tracks circles and lines detected using edge detection, RANSAC, and Hough transforms. Although the LiDAR compass shares some of the same principles as VCs, it can more easily be used as an absolute heading sensor in semi-structured environments (due to the ease of producing DMs). Additionally, the LiDAR compass is generally simple to implement and extremely lightweight, and may be better suited for computationally-restrictive platforms.

II. BACKGROUND

A. Direction Maps

Direction maps (DMs) represent the directions of extractable lines segmented from 2D laser scans of the environment. Lines are extractable if the surfaces generating them are observable by one or more 2D laser scanners. A DM is a set of elements

$$\mathbf{\Phi} = (\phi^{(1)}, \ldots, \phi^{(N)}), \tag{1}$$

where $\phi^{(i)} \in [0, \pi),\ \ i = 1, \ldots, N$, represent the directions of the normal vectors of lines in the environment. The absence of positional information about the lines means that a many-to-one relationship between lines extracted from the environment and line directions in the map is possible. For example, parallel walls in a hallway are represented by one entry in a DM. This also facilitates mounting the laser scanner on the robot because only its orientation relative to the robot coordinate frame matters.

The LiDAR compass uses one of two types of DMs: a static *a priori* DM $\bar{\mathbf{\Phi}}$ with little or no uncertainty, or a dynamically estimated DM $\hat{\mathbf{\Phi}}$ with local directions that are frequently added or removed. If directions of the dominant surfaces in the environment (e.g., walls, sides of buildings) can be determined *a priori*, then $\bar{\mathbf{\Phi}}$ is used. For example, $\bar{\mathbf{\Phi}} = (0, \frac{\pi}{2})$ is likely sufficient for an indoor environment. However, if the environment does not have clearly dominant surfaces and a DM cannot be determined *a priori* (e.g., many diagonal hallways that do not follow the general layout of a building, areas populated by furniture or other objects with unknown orientations), $\hat{\mathbf{\Phi}}$ is used. The direction estimate and uncertainty of these locally added directions are updated as they continue to be observed. Local directions that are no longer being observed are removed from $\hat{\mathbf{\Phi}}$. This maintains a low-dimensional DM, which lowers the risk of erroneous data association and improves robustness to outliers (e.g., lines extracted from noise).

Because the entries of $\hat{\mathbf{\Phi}}$ are observed from a robot heading $\theta \in [-\pi, \pi)$, which itself is uncertain, the entries of $\hat{\mathbf{\Phi}}$ are correlated to θ and to each other. As a result, the LiDAR compass actually estimates the joint heading/DM state $\mathbf{\Theta} = (\theta, \hat{\mathbf{\Phi}})$ and its dense covariance matrix $\mathbf{P}$. The details of adding and removing entries or updating $\hat{\mathbf{\Phi}}$ are described in Section III-C. However, in the common situation when an *a priori* DM is available, $\bar{\mathbf{\Phi}}$ is used in place of $\mathbf{\Theta}$ and is considered an absolute reference with which θ can be estimated.

B. Line Extraction

Line extraction algorithms can generally be divided into two steps: line segmentation and line fitting. Line segmentation determines the groups of points that form lines, and line fitting determines the coefficients of the equation describing each line and their uncertainties. A thorough comparison of line segmentation algorithms for 2D laser scanners was performed by Nguyen *et al*. [13]. They concluded that split-and-merge techniques are the preferred choice for real-time SLAM applications. After removing outlying points based on the proximity to their neighbours, split-and-merge is used for line segmentation in this study. Line fitting is performed using the weighted least squares approach described by Pfister *et al*. [14].

The input to the line extraction algorithm is the range and bearing data (and their uncertainties) generated by a 2D laser scanner. However, unlike most applications of line extraction, the output is simply the estimated directions $z^{(j)}$ of the extracted lines expressed in the LiDAR frame, as well as their corresponding scalar variances $\sigma^2_{z^{(j)}}$; i.e.,

$$\mathbf{Z} = (z^{(1)}, \ldots, z^{(M)}), \tag{2}$$

where $z^{(j)} \in [0, \pi),\ \ j = 1, \ldots, M$ and $z^{(j)} \sim \mathcal{N}(0, \sigma^2_{z^{(j)}})$.

C. Observation Model

The observation model produces predicted line directions in the LiDAR frame given the DM and the current heading of the robot. The predicted observation of the i^{th} map entry at time k in the LiDAR frame is then

$$\hat{z}^{(i)}_k = \phi^{(i)} - (\theta_k + \beta), \tag{3}$$

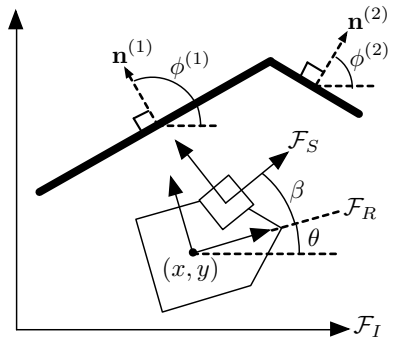

Fig. 1. An illustration of the robot heading, entries of a DM, and the inertial ($\mathcal{F}_I$), robot ($\mathcal{F}_R$), and sensor ($\mathcal{F}_S$) coordinate frames. The entries of a DM are the directions of their normal vectors $\mathbf{n}^{(i)}$. Each predicted observation $\hat{z}_k^{(i)}$ is a DM entry $\phi^{(i)}$ expressed in $\mathcal{F}_S$.

where β is the (constant) orientation of the sensor in the robot coordinate frame. The map directions, sensor orientation, robot heading, and coordinate frames are illustrated in Figure 1. To simplify, $\beta = 0$ will be assumed for the remainder of this paper. Note that the observation space is independent of the positions of the robot, laser scanner (relative to the robot), and lines in the environment.

III. LIDAR COMPASS

The LiDAR compass provides heading estimates to a vehicle navigating within an environment containing line-extractable surfaces. These surfaces could be walls (indoors and outdoors), furniture, cars, etc. For each laser scan, line extraction is performed (as described in Section II-B), followed by a heading estimate by the LiDAR compass.

A. Heading Estimation with an A Priori DM

When an *a priori* DM $\bar{\mathbf{\Phi}}$ is available, an estimate of the robot heading θ_{k+1} is calculated given $\bar{\mathbf{\Phi}}$, θ_k, and the observed line directions $\mathbf{Z}_k$. This requires establishing correspondences among $\bar{\mathbf{\Phi}}$ and $\mathbf{Z}_k$. First, a set of predicted observations $\hat{z}^{(1)}, \ldots, \hat{z}^{(N)}$ are generated from (3); one for each entry in $\bar{\mathbf{\Phi}}$. Next, the total variance $s^{(j)}$ in the sensor frame of each observation $z^{(j)}$ is calculated by combining the uncertainty of the observation and the robot heading ($\sigma_{\theta_k}^2 + \sigma_{z^{(j)}}^2$). Data association is performed by selecting the predicted observation $\hat{z}^{(i)}$ with the minimum Mahalanobis distance $r^{(j,i)}$ from the observation $z^{(j)}$; i.e.,

$$r^{(j,i)} = \min\left(\frac{[z^{(j)} - \hat{z}^{(i)}]^2}{s^{(j)}}, \frac{[z^{(j)} - (\hat{z}^{(i)} + \pi)]^2}{s^{(j)}}\right), \quad (4)$$

$$i^* = \underset{i}{\operatorname{argmin}}\ r^{(j,i)}, \ \text{ for } i = 1, \ldots, N, \quad (5)$$

where i^* is the index of $\bar{\mathbf{\Phi}}$ that most likely corresponds with the j^{th} observation. Note that smallest of the two terms in (4) is taken to account for the bounds on entries in DM; each predicted observation has an equivalent prediction on the opposite side of the unit circle. If the Mahalanobis distance $r^{(j,i^*)}$ is below a threshold, the j^{th} observation is used to update the robot heading much like the observation update of a Kalman filter; i.e.,

$$\theta_{k+1} = \theta_k + \frac{\sigma_{\theta_k}^2}{s^{(j)}}\left(z^{(j)} - \hat{z}^{(i^*)}\right), \quad (6)$$

$$\sigma_{\theta_{k+1}}^2 = \left(1 - \frac{\sigma_{\theta_k}^2}{s^{(j)}}\right)\sigma_{\theta_k}^2, \quad (7)$$

where $\hat{z}^{(i^*)}$ may require a rotation by π, depending on the result of (4).

B. Heading Estimation with a Local DM

When no *a priori* DM is available, an estimate of the robot heading θ_{k+1} is calculated given the previous joint heading/DM state $\mathbf{\Theta}_k$ and the observed line directions $\mathbf{Z}_k$. As before, the predicted observations $\hat{z}^{(1)}, \ldots, \hat{z}^{(N)}$ are calculated using (3). However, because the uncertainties of the local map entries are correlated with θ, the total variance of each of the predicted observations is now

$$s^{(j,i)} = \mathbf{C}^{(i)}\mathbf{P}_k^-[\mathbf{C}^{(i)}]^T + \sigma_{z^{(j)}}^2, \quad (8)$$

where $\mathbf{C}^{(i)} = \begin{bmatrix}-1 & 0^{(1)} & \ldots & 1^{(i)} & \ldots & 0^{(N)}\end{bmatrix}$ and $\mathbf{P}_k$ is the covariance matrix of $\mathbf{\Theta}_k$. Once again, the Mahalanobis distance determines the maximum likelihood correspondences among $\mathbf{Z}_k$ and $\hat{\mathbf{\Phi}}_k$, and is calculated using (4) and (5) as before with $s^{(j,i)}$ substituted for $s^{(j)}$. However, unassociated observations are now used to add new entries to $\hat{\mathbf{\Phi}}$ with the inverse observation model; i.e.,

$$\mathbf{F} = \begin{bmatrix}1 & 0^{(1)} & \ldots & 0^{(N)}\end{bmatrix}, \quad (9)$$

$$\mathbf{\Theta}_{k+1} = \begin{bmatrix}\theta_k \\ z_k^{(j)} + \theta_k\end{bmatrix}, \quad (10)$$

$$\mathbf{P}_{k+1} = \begin{bmatrix}\mathbf{P}_k & \mathbf{P}_k\mathbf{F}^T \\ \mathbf{F}\mathbf{P}_k & \sigma_{\theta_k}^2 + \sigma_{z^{(j)}}^2\end{bmatrix}. \quad (11)$$

Associated observations update the heading estimate (and the local DM) in a fashion similar to (6) and (7); i.e.,

$$\mathbf{K}^{(j)} = \frac{\mathbf{P}_k[\mathbf{C}^{(i^*)}]^T}{s^{(j,i^*)}}, \quad (12)$$

$$\mathbf{\Theta}_{k+1} = \mathbf{\Theta}_k + \mathbf{K}^{(j)}\left(z^{(j)} - \hat{z}^{(j)}\right), \quad (13)$$

$$\mathbf{P}_{k+1} = \left(\mathbf{I} - \mathbf{K}^{(j)}\mathbf{C}^{(i^*)}\right)\mathbf{P}_k. \quad (14)$$

Finally, it is possible that the estimated directions of two entries in the local DM converge as they are updated, meaning they represent the same line direction in the environment. These occurrences are detected by calculating the Mahalanobis distance $r^{(m,n)}$ between the expected distance (which is zero) and the estimated distance between each pair (m, n) of map entries; i.e.,

$$\begin{aligned}\mathbf{C}^{(m,n)} &= \left[0, 0^{(1)}, \ldots, 1^{(m)}, \ldots, -1^{(n)}, \ldots, 0^{(N)}\right], \\ s^{(m,n)} &= \mathbf{C}^{(m,n)}\mathbf{P}_{k+1}[\mathbf{C}^{(m,n)}]^T, \\ r^{(m,n)} &= \frac{\left[0 - \left(\phi^{(m)} - \phi^{(n)}\right)\right]^2}{s^{(m,n)}}.\end{aligned} \quad (15)$$

If $r^{(m,n)} < r_{\text{thresh}}$, entries m and n must be merged. First, $\boldsymbol{\Theta}$ and $\mathbf{P}$ are updated using (12)-(14) with $z^{(j)} = 0$, $\hat{z}^{(j)} = \left(\phi^{(m)} - \phi^{(n)}\right)$ and $\mathbf{C}^{(m,n)}$ in place of $\mathbf{C}^{(i^*)}$. The n^{th} map entry is then removed from $\boldsymbol{\Theta}$ and $\mathbf{P}$. Note that the succinct size of local DM makes a comparison of each pair of map entries computationally feasible.

C. Brightnesses

When no *a priori* DM is available and a local DM is used, only frequently observed directions are kept in $\hat{\boldsymbol{\Phi}}$, which maintains low-dimensionality and prevents erroneous data associations. To determine which entries make up the local DM, each entry is assigned a *brightness* $b \in \mathbb{N}$. Observing map entries causes them to *brighten*, while unobserved map entries *fade*. When a new entry is added to the DM, it is assigned an initial brightness $b_0 > 0$. For each observation in $\mathbf{Z}_k$ that is associated with $\phi^{(i)}$, $b^{(i)}$ increases (brightens) by one. If $\phi^{(i)}$ is not observed at time k, $b^{(i)}$ decreases (fades) by one. Map entries can brighten up to the maximum brightness B or fade to zero. Map entries fading to zero are removed from the DM, resulting in a succinct DM that is comprised solely of map entries actively being observed.

By tuning b_0 and B in conjunction with the frequency of the heading updates, one can specify the amount of time it takes for map entries to reach maximum brightness or fade to zero. Erroneously-added entries (such as a lines extracted from noise) are quickly removed from the map because they tend not to be re-observed. Additionally, the brightness of map entries are used to adjust the magnitude of an heading update by inserting a scaling factor to the gain calculation, e.g., changing (12) to

$$\mathbf{K}^{(j)} = \frac{b^{(i^*)}}{B} \frac{\mathbf{P}_k [\mathbf{C}^{(i^*)}]^T}{s^{(j,i^*)}}. \tag{16}$$

This results in heading updates associated with bright map entries (e.g., the direction of the walls of a room) having a greater effect on the heading estimate than faded map entries.

IV. LOCALIZATION WITH A LIDAR COMPASS

This section describes how the LiDAR compass algorithms described in Section III could be used as part of a localization algorithm for a mobile robot or unmanned ground vehicle. When used with an *a priori* DM (Section III-A), this is analogous to providing the robot with an absolute heading sensor (e.g., compass, sun sensor). However, it can be used in any environment with line-extractable surfaces. When used without an *a priori* DM (Section III-B), the LiDAR compass is analogous to a gyroscope, in that it provides dead-reckoning heading estimation. However, unlike a gyroscope, it estimates the heading directly (without integration). In both cases, the LiDAR compass provides a correction to the robot pose $\mathbf{q} = (x, y, \theta) \in SE(2)$.

It is assumed that the robot performs a motion update that provides an *a priori* estimate $\mathbf{q}_{k+1}^-$ (and its covariance matrix $\mathbf{P}_{k+1}^-$). For example, a commonly used motion model for a differential drive vehicle that provides such an estimate is described by Thrun *et al*. [15, Chapter 5]. Wheel encoders or motion commands are generally used to calculate the *a priori* estimate. The LiDAR compass then calculates the *a posteriori* estimate with slight modifications to the equations given in Section III. When using an *a priori* DM, (6) and (7) become

$$\mathbf{K}^{(j)} = \frac{\mathbf{P}_{k+1}^- \begin{bmatrix} 0 & 0 & -1 \end{bmatrix}^T}{s^{(j)}}, \tag{17}$$

$$\mathbf{q}_{k+1}^+ = \mathbf{q}_{k+1}^- + \mathbf{K}^{(j)} \left(z^{(j)} - \hat{z}^{(i^*)} \right), \tag{18}$$

$$\mathbf{P}_{k+1}^+ = \left(\mathbf{I} - \mathbf{K}^{(j)} \begin{bmatrix} 0 & 0 & -1 \end{bmatrix} \right) \mathbf{P}_{k+1}^-, \tag{19}$$

to accommodate updates to the full robot pose. When the LiDAR compass is used with local maps only, the joint heading/DM state $\boldsymbol{\Theta}$ is augmented to include the full pose of the robot, and becomes $\boldsymbol{\Theta} = (\mathbf{q}, \hat{\boldsymbol{\Phi}})$. The update equations (12)–(14) are used with this new $\boldsymbol{\Theta}$ and

$$\mathbf{C}^{(i)} = \begin{bmatrix} 0 & 0 & -1 & 0^{(1)} & \dots & 1^{(i^*)} & \dots & 0^{(N)} \end{bmatrix}.$$

V. EXPERIMENTAL RESULTS

Experiments were performed using a Clearpath Robotics Husky A200 mobile robot equipped with a SICK LMS111 2D laser scanner oriented to scan in the horizontal plane. The laser scanner has a field of view of $270°$, a resolution of $0.5°$, and a range of 20 m. The test environment was at Queen's Innovation Park in Kingston, Ontario, Canada, where the robot was driven through a warehouse-like area with a concrete floor measuring approximately 40 m $\times$ 60 m (illustrated in Fig. 2). Various shelves, equipment, and (moving) people were present in the environment during the experiments. The robot was driven at approximately 0.5 m/s and collected laser scans at 50 Hz. The primary goal of the experiments was to demonstrate the effectiveness of the primary focus of the LiDAR compass: heading estimation. However, a simple motion model using wheel encoder data (10 Hz) was also implemented to test using the LiDAR compass as part of a localization algorithm (as described in Section IV).

A. Ground Truth

Ground truth poses at 15 locations were determined by manually measuring the relative distance and angle between locations in the environment. These measurements formed a nonlinear graph optimization problem describing the locations of a set of points (the graph is shown in Fig. 2). By considering the uncertainty of the manual measurements, the optimal estimate of the ground truth poses was solved in MATLAB®by using the Levenberg-Marquardt algorithm. The heading at each ground truth pose was directly measured with respect to known surfaces (e.g., walls); therefore, its uncertainty is constant. The resulting ground truth poses are shown in Appendix VI. The uncertainty in the ground truth is accounted for in the error analysis.

B. Results and Error Analysis

Although several experiments were performed, this paper reports on five trials in which the robot was manually driven between the ground truth poses in the sequence

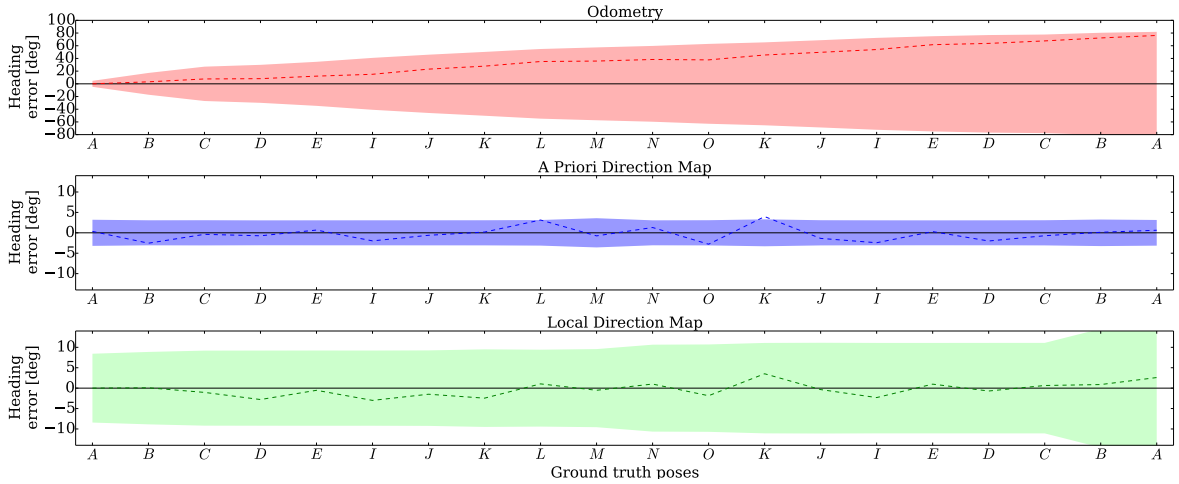

Fig. 4. The heading errors using odometry (top), a LiDAR compass with an *a priori* DM (middle), and a LiDAR compass with a local DM (bottom). Note the scale difference in the odometry plot. The 3σ uncertainty in the errors includes both the uncertainty in the heading estimate and the uncertainty in the ground truth. The RMS errors are 43.67° (odometry), 1.74° (*a priori DM*), and 1.72° (local DM).

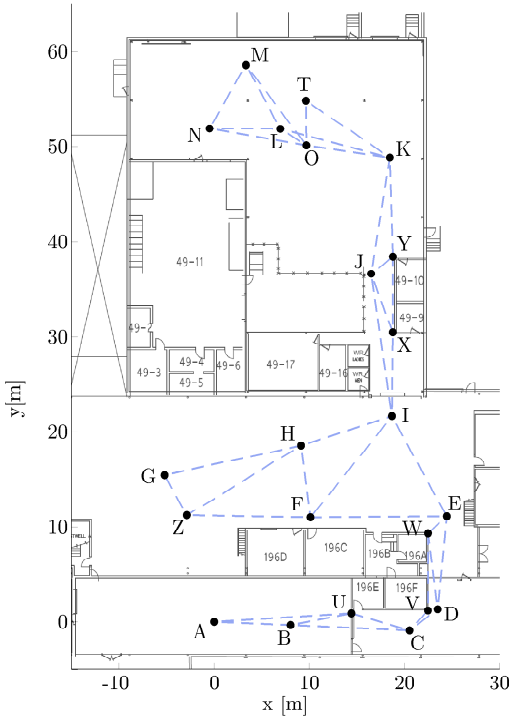

Fig. 2. A blueprint of Innovation Park showing the graph of manually measured positions that was optimized for use as ground truth (see Table I). Poses A–O were marked on the ground to compare with the localization algorithms. A dashed line indicates that a range and bearing measurement was taken between the connected nodes.

$(A, B, C, D, E, I, J, K, L, M, N, O, K, J, I, E, D, C, B, A)$, not necessarily driving in straight lines between poses. The total travel distance of this sequence is over 200 m. The robot was briefly parked at the ground truth poses to record the heading estimated by the LiDAR compass and pose estimates of the localization algorithms. Although the manual parking was performed carefully, it nevertheless contributed a small amount of error to the ground truth comparisons. Three different heading estimation methods were performed simultaneously: a LiDAR compass with an *a priori* DM $\bar{\mathbf{\Phi}} = (0, \frac{\pi}{2})$, a LiDAR compass with a local DM (initially empty), and odometry (estimating the heading using the difference of the left and right wheel encoders). The full localization form of these algorithms was also performed using a wheel encoder based motion update. All five trials produced similar results, and the resulting path estimates by the three algorithms (when used for localization) in one of the trials is shown in Fig. 3. The errors in the heading estimates at each of the ground truth poses are illustrated in Fig. 4 with 3σ uncertainty. Assuming the heading estimates and ground truth measurements are independent Gaussian random variables, the variance of the difference between them (i.e., the uncertainty associated with the error) is the sum of the independent covariances.

C. Discussion

The results presented in Figs. 3 and 4 demonstrate the effectiveness of using a LiDAR compass for robot heading estimation. There was essentially no heading error growth over the course of the trial. Although bounded error growth was expected with an *a priori* DM (because it has an absolute reference), there was very little error when only a local DM was used as well. The performance of the local DM is likely linked to the environment, which consistently had similarly directed surfaces in most areas. As a result, these directions were consistently part of the local DM. Future work will test the effectiveness of using the local DM in a more challenging environment with a larger variety of directions.

The *a priori* DM used by the LiDAR compass was determined simply by observing the directions of dominant surfaces of the environment and giving that information to the robot (i.e., it requires little to no work). As a result, it is possible that in the environments suitable for a LiDAR compass, providing a *a priori* DM (however small) may always be possible. To maintain the advantages of both using

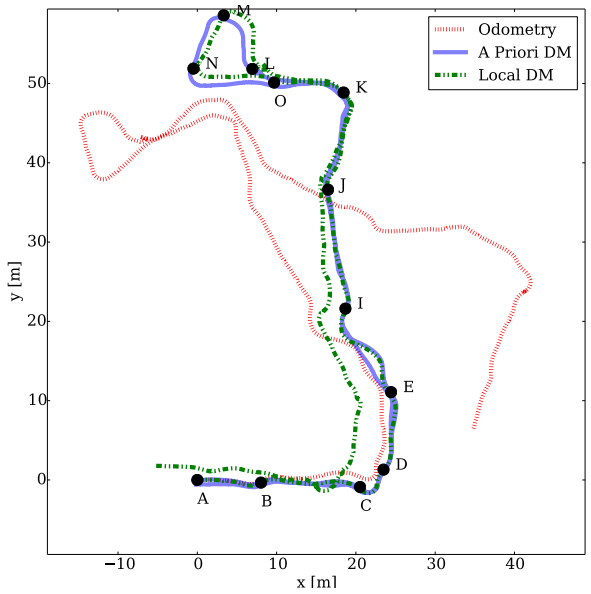

Fig. 3. The robot's path over the course of a single trial, as estimated by the localization algorithms using a LiDAR compass with an *a priori* DM, a LiDAR compass with a local DM, and odometry. The visited ground truth poses are shown for reference. The total path length was over 200 m.

an *a priori* DM (bounded error growth) and a local DM (flexibility in differently directed environments), future work will consider a hybrid approach, where the LiDAR compass has two DMs. Observations will first be compared to a static, *a priori* DM, and will then be passed onto a local DM if no associations are found. This way, the local DM will "take over" if the robot observes no surfaces in its *a priori* DM.

VI. CONCLUSIONS

This paper introduces an extremely lightweight method to effectively bound the growth of heading errors for a mobile robot in a semi-structured environment. It takes advantage of reoccurring surface directions to greatly reduce or bound the error growth from dead reckoning. The main contribution of this algorithm is the simplicity of its heading updates and map representation, for both computation and ease of implementation. In environments where absolute heading sensors such as a compass, magnetometer, or sun sensor are generally inaccurate or ineffective (e.g., most indoor environments), the LiDAR compass offers an absolute reference if line-extractable surfaces are available. Future work will include a hybrid *a priori*/local DM implementation, testing outdoors, and comparisons with other approaches that also attempt to reduce heading errors introduced by odometry.

APPENDIX

Table I shows the computed ground truth poses.

ACKNOWLEDGMENT

The authors would like to thank Oscar Rielo, Perry Ross, and the staff at Innovation Park for help with experiments.

TABLE I

THE LOCATION AND UNCERTAINTY OF THE GROUND TRUTH POSES.

Pose	x [m]	y [m]	θ [deg]	σ_x [m]	σ_y [m]	σ_θ [deg]
A	0.000	0.000	0.0	0.000	0.000	0.5
B	8.026	-0.336	22.2	0.002	0.037	0.5
C	20.530	-0.885	-31.9	0.004	0.080	0.5
D	23.483	1.304	57.6	0.011	0.081	0.5
E	24.442	11.079	124.7	0.041	0.082	0.5
F	10.114	10.978	160.0	0.041	0.092	0.5
G	-5.210	15.42	66.1	0.043	0.119	0.5
H	9.126	18.561	27.1	0.046	0.093	0.5
I	18.679	21.609	63.8	0.051	0.083	0.5
J	16.492	36.610	103.0	0.072	0.084	0.5
K	18.453	48.861	136.4	0.104	0.083	0.5
L	6.944	51.843	128.6	0.104	0.090	0.5
M	3.333	58.599	163.1	0.107	0.094	0.5
N	-0.492	51.877	-111.1	0.104	0.099	0.5
O	9.686	50.117	-9.1	0.104	0.088	0.5

REFERENCES

[1] G. Grisetti, R. Kummerle, C. Stachniss, and W. Burgard, "A Tutorial on Graph-Based SLAM," *IEEE Intelligent Transportation Systems Magazine*, vol. 2, no. 4, pp. 31–43, 2010.

[2] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental Smoothing and Mapping," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1365–1378, Dec. 2008.

[3] I. J. Cox, "Blanche- An experiment in guidance and navigation of an autonomous robot vehicle," *IEEE J. Robot. Automat.*, vol. 7, no. 2, pp. 193–204, Apr. 1991.

[4] J. Vandorpe, H. Van Brussel, and H. Xu, "Exact dynamic map building for a mobile robot using geometrical primitives produced by a 2D range finder," in *Proc. IEEE International Conference on Robotics and Automation*, Minneapolis, USA, Apr. 1996, pp. 901–908.

[5] A. Garulli, A. Giannitrapani, A. Rossi, and A. Vicino, "Mobile robot SLAM for line-based environment representation," in *Proc. IEEE Conference on Decision and Control*, Seville, Spain, Dec. 2005, pp. 2041–2046.

[6] V. Nguyen, A. Harati, A. Martinelli, and R. Siegwart, "Orthogonal SLAM: a step toward lightweight indoor autonomous navigation," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, Oct. 2006, pp. 5007–5012.

[7] B.-W. Kuo, H.-H. Chang, Y.-C. Chen, and S.-Y. Huang, "A light-and-fast SLAM algorithm for robots in indoor environments using line segment map," *Journal of Robotics*, vol. 2011, 2011.

[8] F. Labrosse, "The visual compass: Performance and limitations of an appearance-based method," *Journal of Field Robotics*, vol. 23, no. 10, pp. 913–941, 2006.

[9] N. Bellotto, K. Burn, E. Fletcher, and S. Wermter, "Appearance-based localization for mobile robots using digital zoom and visual compass," *Robotics and Autonomous Systems*, vol. 56, no. 2, pp. 143–156, Feb. 2008.

[10] J. Košecká and W. Zhang, "Video Compass," in *Proc. 7th European Conf. Comp. Vis.*, 2002, pp. 476–491.

[11] J. Montiel and A. Davison, "A visual compass based on SLAM," in *IEEE International Conference on Robotics and Automation*, 2006, pp. 1917–1922.

[12] G. L. Mariottini, S. Scheggi, F. Morbidi, and D. Prattichizzo, "An accurate and robust visual-compass algorithm for robot-mounted omnidirectional cameras," *Robotics and Autonomous Systems*, vol. 60, no. 9, pp. 1179–1190, Sep. 2012.

[13] V. Nguyen, A. Martinelli, N. Tomatis, and R. Siegwart, "A comparison of line extraction algorithms using 2D laser rangefinder for indoor mobile robotics," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, Edmonton, Canada, Aug. 2005, pp. 1929–1934.

[14] S. T. Pfister, S. I. Roumeliotis, and J. W. Burdick, "Weighted line fitting algorithms for mobile robot map building and efficient data representation," in *Proc. IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, Sep. 2003, pp. 14–19.

[15] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2006.