Robust Pose-Graph SLAM Using Absolute Orientation Sensing

Saurav Agarwal, Karthikeya S. Parunandi, and Suman Chakravorty 🔎

Abstract—It is known that in the simultaneous localization and mapping (SLAM) problem when a robot's orientation is known, an estimation of the history of its poses can be formulated as a standard linear least squares problem. In this letter, we exploit this property of SLAM to develop a robust pose-graph SLAM framework that uses absolute orientation sensing. Our contribution are as follows: 1) we show that absolute orientation can be estimated using local structural cues; and 2) we develop a method to incorporate absolute orientation measurements in both the front and back-end of pose-graph SLAM. We also demonstrate our approach through extensive simulations and a physical real-world demonstration along with comparisons against existing state-of-the-art solvers that do not use absolute orientation.

Index Terms—SLAM, localization, mapping, range sensing, visual-based navigation.

I. INTRODUCTION

T IS well understood that a difficult problem in mobile robot autonomy is the challenge of long-term navigation using only on-board sensors in large GPS-denied environments. An important practical example are material handling robots that move goods (boxes, pallets etc.) in large warehouses and distribution centers. Since installing beacons, markers or guide cables is expensive, robust, accurate localization and mapping is necessary for reliable operation. A common technique to solve the SLAM problem is to use relative pose-graph SLAM. Pose-graph SLAM uses a two-pronged approach; (i) a front-end which maintains an estimate of the robot pose using correlative scan matching [1] or other suitable techniques and computes data association between current and past observations, and (ii) a back-end which solves the non-linear optimization to compute the history of robot poses.

A. Problems With Existing SLAM Solutions

A major caveat of existing SLAM techniques is that to correct estimation drift they often rely on loop closure, i.e., revisiting previously seen locations and correctly associating sensor information to data previously stored in the map. There are two key problems that must be noted; (i) loop closure is *sensitive to data*

Manuscript received September 10, 2018; accepted December 28, 2018. Date of publication January 16, 2019; date of current version February 8, 2019. This letter was recommended for publication by Associate Editor U. Frese and Editor C. Stachniss upon evaluation of the reviewers' comments. This work was funded in part by NSF ECCS grant 1637889 and I-Corps grant 1740544. (*Corresponding author: Suman Chakravorty.*)

The authors are with the Department of Aerospace Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: saurav6@gmail.com; s.parunandi@tamu.edu; schakrav@tamu.edu).

Digital Object Identifier 10.1109/LRA.2019.2893436

association accuracy, thus wrong data association can lead to catastrophic failure of the SLAM system as shown in Fig. 1(a), and (ii) data association reliability is limited by localization accuracy, thus localization drift causes map quality to degrade as the scale of environment increases.

B. Contributions

We propose a method for robust 2D SLAM that fuses absolute orientation sensing (using cameras that track stable structural features), with range-scan measurements using a LiDAR. Realistic simulation studies show that **our proposed method does not fail in mapping whereas existing state-of-the-art methods fail** $\approx 40\% - 50\%$ of the times (see Table I). We conduct a detailed analysis of the effect of noisy relative orientation measurements and show that absolute orientation measurements are critical to achieving robust localization and mapping. We demonstrate our approach successfully on a physical system in a real-world setting with a commercially available mobile robot platform.

Furthermore, through physical experiments we show that absolute orientation information may be robustly sensed in indoor scenarios by leveraging structural features. For example, in industrial buildings, the ceiling lighting or corrugation is usually aligned along one direction. The benefit of this approach can be clearly seen in Fig. 1.

C. Limitations

For 3D applications, vertical columns or pillars may be used to infer roll and pitch information however we limit the scope of this work to 2D planar SLAM. Further, the current physical application of this work is limited to environments with ceilings that have a Manhattan world structure (grid pattern) which presents strong structural cues for heading (yaw) information. For instance, this assumption is satisfied by open ceiling structures like warehouses, or office buildings that are laid out in rectangular grids. However, such global orientation information may not be available in environments like airports which do not have a fixed direction, might be circular etc. Nonetheless, we hope that this work recognizes the central importance of global orientation in SLAM and leads to work that can get such global orientation information, in a robust fashion, in a wide variety of environments.

II. RELATED WORK

The first non-linear optimization based approach to solving the full SLAM problem was introduced by the seminal work of [4]. Several other works [5]–[8] made important contributions to the original formulation. The works of [9]–[13] have sought

2377-3766 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.



(a) Catastrophic mapping failure results in a twisted map that is unusable by a robot for real-world tasks. Such failures are quite common with SLAM implementations.



(b) Reliable mapping that is suitable to be used by a robot for real-world tasks.

Fig. 1. Existing state-of-the-art vs. proposed approach: Figure (a) shows an example of mapping failure for a building with a floor area of 78,240 sq. ft. The floor plan is based on an HEB grocery store located in College Station, Texas. We used a state-of-the-art front-end [1], [2] and g2o [3] for the back-end. The solution failed in $\approx 40\% - 50\%$ of the experiments that were run. Figure (b) shows an example of successful mapping for the same environment. We used our novel approach which fuses absolute orientation information with a state of the art front-end [1], [2] and our back-end graph solver. This approach succeeded in every one of the experiments that were run.

TABLE I Comparison of Existing State-of-the-Art vs. Our Approach. Both Methods Use the Same Scan Matching Approach. Our Approach Suffered From 0 Catastrophic Failures

Front-end	Back-end	Lidar Range	Total Runs	Failures	Average RMSE
Correlative Scan Matching	g2o	30	50	21	3.90m
Our Approach	LOGO	30	50	0	0.36m
Correlative Scan Matching	g2o	20	50	25	5.60m
Our Approach	LOGO	20	50	0	0.35m

to exploit structural properties of SLAM in order to decouple the non-linearities induced by orientation. In [9] the authors develop an incremental SLAM approach that exploits orientation and position separation in 2D range-scan mapping. The authors in [14] present a general on-manifold optimization based approach for the estimation of orientation from noisy relative measurements that are corrupted by outliers. The reader is directed to [15] for a recent survey on 3D rotation estimation methods. In [10], the authors present the Linear Approximation for pose Graph Optimization (LAGO) method for planar relative posegraph SLAM that decouples robot orientation and position estimation into two successive linear problems with a key advantage of reduced risk of convergence to local minima while also providing a robust initial guess for iterative optimization. In [11] the authors develop MOLE2D, a multi-hypothesis approach to global orientation estimation from relative measurements that does not suffer from local minima even in the case of high noise. In [13] a method is developed for 2D feature-based SLAM that separates orientation and position estimation by leveraging relative feature measurements. In relation to these works, our current work leverages orientation cues available in the environment to sense absolute orientation and fuses that information in the front-end and back-end of a pose-graph SLAM formulation. This work goes a step further in that, instead of estimating orientation from relative measurements, it leverages absolute orientation sensing to solve a linearized least squares problem in the back-end while making the SLAM front-end highly robust to scan registration errors.

In [16] the authors present an odometry technique that uses a ceiling facing RGB-D camera (e.g., Microsoft Kinect). The method of [16] does not disambiguate the principal direction from perpendicular ceiling lines. Only changes in the principal direction are used to aid in the ICP-based odometry, thus the system is still prone to drift but performs better than dead reckoning and wheel odometry. The closest to our approach is the Compass-SLAM approach of [17] that utilizes a compass to get global orientation within an early pose Graph SLAM framework. However, our orientation sensing system is vision based, and thus, not prone to the errors typical to a compass. We also analytically show the compounding effect of relative orientation errors in a SLAM solution and show the relative advantage with respect to GraphSLAM approaches without orientation information.

III. PRELIMINARIES AND PROBLEM

Let $\mathbf{x} = {\mathbf{x}_0, ..., \mathbf{x}_n}$ be a set of n + 1 poses, describing the robot position and orientation at each time k. In 2D (planar) problems, $\mathbf{x}_k = [\mathbf{p}_k^T \theta_k]^T \in SE(2)$, where $\mathbf{p}_k \in \mathbb{R}^2$ is the position and $\theta_k \in SO(2)$ is the heading. Let ξ_{ij} be a relative measurement of pose j w.r.t pose i then,

$$\xi_{ij} = \begin{bmatrix} {}^{l} \boldsymbol{\Delta}_{ij} = \mathbf{R}_{i} (\mathbf{p}_{j} - \mathbf{p}_{i}) \\ \delta \theta_{ij} = \theta_{j} - \theta_{i} \end{bmatrix}$$
(1)

where \mathbf{R}_i is the rotation matrix composed by θ_i . In the general setting, ξ_{ij} is corrupted by noise, thus $\hat{\xi}_{ij} = \xi_{ij} + v_{ij}$, where v_{ij} is assumed to be zero-mean Gaussian. Let ${}^{l}\boldsymbol{\Delta}$ be the vector of relative position measurements in the local frame at each pose. If robot orientation θ^* is known at each pose, then the SLAM problem simply becomes

$$^{w}\mathbf{\Delta} = \mathbf{R}(\boldsymbol{\theta}^{*})^{\mathrm{T}\ l}\mathbf{\Delta}.$$
 (2)

Algorithm 1: Orientation Estimation from Structural Cues.

1 Input: I, W, b

- 2 I is the input image
- 3 $l \leftarrow$ Extract line features from image I
- 4 $[\theta_1, \theta_2, \dots, \theta_n] \leftarrow$ Compute orientation of the lines in local frame of image *I*
- 5 $h \leftarrow$ Create histogram of the orientation data with bins of width b in range $[0, 2\pi)$
- 6 $i_{max} \leftarrow$ Get index of bin with maximum measurements
- 7 $h_W \leftarrow$ Extract a window of width W around the bin i_{max}
- $\begin{array}{l} i_{max} \\ 8 \ \theta_c, \sigma_{\theta_c}^2 \leftarrow \text{Compute the weighted mean of observations} \\ \text{in window} \end{array}$
- 9 $\theta_c \leftarrow \text{Wrap } \theta_c \text{ in range } [0, \pi)$
- 10 return $\theta_c, \sigma_{\theta}^2$;

We know that ${}^{w}\Delta = \mathbf{A}'\mathbf{p}$ where \mathbf{A}' is a matrix composed of elements in the set $\{-1, 0, 1\}$ and \mathbf{p} is the vector of robot positions in the global frame. Thus, when robot orientation $\boldsymbol{\theta}$ is known, the position estimation problem is linear. Moreover, when unbiased global heading measurements are available, the problem can be very accurately linearized. This is the key insight that provides our method with high localization accuracy providing unparalleled robustness in the front-end and enabling *computationally low-cost linear-least squares solution* for the back-end.

IV. METHOD

We now proceed to describe our approach in detail. Our method comprises three key aspects:

- 1) Using commonly occurring structural cues to sense absolute orientation of the robot. In an indoor setting, absolute orientation is measured relative to the building North.
- 2) Fusing absolute orientation measurements to the SLAM front-end, i.e., in our case to a scan matching algorithm.
- Solving a batch optimization problem, to compute global estimates at loop closure by fusing relative pose measurements and absolute orientation sensing.

The next section deals with the issue of orientation sensing. In Section IV-B we describe how the heading measurements are fused with the front-end in a filtering based scheme. Section IV-C describes the batch optimization method to solve for loop closure and finally in Section IV-D a discussion is presented for our approach.

A. Absolute Orientation Sensing

We seek independent absolute orientation estimates of the robot heading. In an indoor environment, we may use the relative heading of the robot w.r.t the building's fixed North. Our orientation sensing method works by detecting structural features of the environment. In indoor environments such as offices, factories, warehouses etc. the ceiling structure usually has straight line features that are easy to detect. For example, ceiling corrugation in most industrial buildings is aligned along one direction which can be detected by a ceiling facing camera. Lines in the horizontal group may be used to provide heading information and vertical lines are quite reliable for deducing roll and pitch. In this work, we restrict ourselves to tracking horizontal



Fig. 2. The LOGO-SLAM architecture. A flowchart depicting how sensor data flows and the various computation modules.

line features to determine heading. Algorithm 1 describes the process of estimating ceiling direction.

Figure 6(c) shows a polar plot of the histogram generated in Algorithm 1 in one of our simulation runs. Note that line direction is inherently ambiguous, i.e., it may not be possible to differentiate North from South (likewise from East to West). Therefore gyro data is used in the intermediate time between absolute orientation measurements to propagate the robot heading. Gyros usually provide data at > 100Hz and therefore can be used to account for the angle wrap-around issue in absolute orientation detection. To estimate the robot heading, initial heading at time t_0 is assumed to be known (usually $\theta_0 = 0$). In the proceeding section, we describe how absolute orientation measurements are fused in the SLAM front-end.

B. Heading Assisted Front-End

We develop a modified version of the correlative scan matching method of [1] for our front-end. To the scan matching based front-end we add a Kalman filter after the scan match step. This Kalman filter fuses relative orientation estimates from scan matching with absolute orientation estimates to track the robot's heading. In the interest of space we omit the Kalman filter equations as they are well understood and straightforward. However we add one additional step at the update stage of the Kalman filter wherein we use an innovation filter as a form of error checking. This is due to the fact that an erroneous scan match solution can produce corrupt yet confident (inconsistent) estimate of relative pose transformation. If passed into the back-end, an inconsistent relative pose estimate can result in egregious localization and mapping errors. Thus at time t_k if the innovation signal $\tilde{y}_k > \psi$, where ψ is a user defined threshold parameter (set to 30° in our experiments), then the scan match is rejected and the robot relies purely on wheel odometry and gyro to update its localization estimate at time t_k .

C. Backend

Our SLAM back-end uses the graph generated by the frontend along with absolute orientation data and solves a two step optimization problem. The first step is the estimation of robot orientation using the absolute orientation and relative orientation measurements followed by a second step in which a linearized least-squares optimization problem is solved for the robot position. We now proceed to describe both these steps in further detail.

1) Orientation Estimation: Robot orientation $\theta \in (-\pi, \pi]$, thus as the robot navigates, the relative orientation measure-

ments do not provide information about the angle wrap around. Let $\delta \theta_{i,j}$ to be the relative orientation measurement from pose x_i to x_j , then $\delta \theta_{i,j} = \phi(\theta_j - \theta_i)$, where ϕ is the module operator such that $\phi(\theta) \in (-\pi, \pi]$. Thus the regularized relative orientation measurement $\delta \overline{\theta}$ is $\overline{\delta \theta}_{i,j} = \theta_j - \theta_i + 2k_{ij}\pi$. Here k_{ij} is the integer ambiguity. In [10], the authors present an integer ambiguity approximation approach which exploits the fact that relative orientation measurements over a cycle in the graph add up to 0. In our approach, since absolute orientation measurements are available, the integer ambiguity can be simply be calculated as

$$k_{ij} = \operatorname{round}((\overline{\delta\theta}_{i,j} - (\bar{\theta}_j - \bar{\theta}_i))/2\pi).$$
(3)

Once the regularization constants are computed, we formulate a linear estimation problem by stacking together the absolute orientation measurement vector $\overline{\theta}$ and regularized relative orientation measurement vector $\overline{\delta\theta}$ as

$$\boldsymbol{\beta} = \begin{bmatrix} \overline{\boldsymbol{\delta}\boldsymbol{\theta}} \\ \overline{\boldsymbol{\theta}} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{B}' \\ \mathbf{I} \end{bmatrix}}_{\mathbf{B}} \boldsymbol{\theta} + \begin{bmatrix} \mathbf{v}_{\boldsymbol{\delta}\boldsymbol{\theta}} \\ \mathbf{v}_{\overline{\boldsymbol{\theta}}} \end{bmatrix}, \tag{4}$$

where $v_{\delta\theta}$ and $v_{\bar{\theta}}$ are zero mean gaussian noise and B' is a matrix with each row containing elements of set $\{-1, 0, +1\}$. It can be solved for the global orientation estimate as

$$\hat{\boldsymbol{\theta}} = (\mathbf{B}^T \mathbf{R}_{\boldsymbol{\beta}}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \mathbf{R}_{\boldsymbol{\beta}}^{-1} \boldsymbol{\beta}, \qquad (5)$$

and the estimate error covariance is $\Sigma_{\theta} = (\mathbf{B}^T \mathbf{R}_{\beta}^{-1} \mathbf{B})^{-1}$.

2) Position Estimation: Once a global orientation estimate $\hat{\theta}$ is computed, we proceed to compute robot position at each pose. From Eq. (1) we know that a relative pose measurement from pose x_i to x_j contains a relative position measurement Δ as

$${}^{l}\boldsymbol{\Delta}_{ij} = \mathbf{R}_{i}(\mathbf{p}_{j} - \mathbf{p}_{i}), \tag{6}$$

where ${}^{l}\Delta_{ij}$ is the displacement measured in the local frame of pose x_i and $\mathbf{p}_i, \mathbf{p}_j$ are the 2D positions. Abusing notation slightly, let ${}^{l}\hat{\Delta} \sim \mathcal{N}({}^{l}\Delta, {}^{l}\mathbf{R}_{\Delta})$ be the vector of all local relative position measurements. The vector of local relative measurements ${}^{l}\hat{\Delta}$ can be transformed to the global frame similar to Eq. (2). Thus replacing \mathbf{R}_i with $\mathbf{R}(\hat{\theta}_i)$, and using Eq. (2), we can formulate the linear estimation problem as

$$^{w}\hat{\Delta} = \hat{\mathbf{R}}^{l}\hat{\Delta} = \mathbf{A}'\mathbf{p} + ^{w}\mathbf{v}_{\Delta}.$$
 (7)

where $\hat{\mathbf{R}} = \mathbf{R}(\hat{\boldsymbol{\theta}})$ is the corresponding composition of rotation matrices parametrized by the estimated heading $\hat{\boldsymbol{\theta}}$, \mathbf{p} is the vector of robot positions, \mathbf{A}' is a matrix with each row containing elements of the set $\{-1, 0, +1\}$ and ${}^{w}\mathbf{v}_{\Delta} \sim \mathcal{N}(\mathbf{0}, {}^{w}\mathbf{R}_{\Delta} = \mathbf{C}^{T l}\mathbf{R}_{\Delta}\mathbf{C})$ is the noise vector.

In Section IV-C1, we have shown that global heading estimates comprise both absolute and relative measurements, thus heading estimates are correlated. Solving Eq. (7) directly would result in an erroneous estimate as the correlations have not been accounted for yet. Now we proceed to describe how to setup the position estimation problem while correctly incorporating the appropriate error covariances similar to the trick employed in [10], [13]. After computing the orientation estimates $\hat{\theta}$ along with the transformed global relative position measurements we stack them to give us a new measurement vector γ . Then we have

$$\gamma = \mathbf{h}_{w}(^{l}\boldsymbol{\Delta},\boldsymbol{\theta}) + \mathbf{v}_{w} = \begin{bmatrix} \hat{\mathbf{R}}^{l}\hat{\boldsymbol{\Delta}}\\ \hat{\boldsymbol{\theta}} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{A}' & \mathbf{0}\\ \mathbf{0} & \mathbf{I} \end{bmatrix}}_{\mathbf{A}} \begin{bmatrix} \mathbf{p}\\ \boldsymbol{\theta} \end{bmatrix} + \begin{bmatrix} ^{w}\mathbf{v}_{\mathbf{\Delta}}\\ \mathbf{v}_{\boldsymbol{\theta}} \end{bmatrix}.$$
(8)

The error covariance \mathbf{R}_{γ} of measurement vector γ is then given by,

$$\mathbf{R}_{\gamma} = \bar{\nabla} \mathbf{h}_{w} \begin{bmatrix} {}^{l} \mathbf{R}_{\Delta} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{\boldsymbol{\theta}} \end{bmatrix} \bar{\nabla}^{T} \mathbf{h}_{w}$$
(9)

where $\nabla \mathbf{h}_w$ is the Jacobian of measurement function \mathbf{h}_w Eq. (8). Finally, the solution to the linear estimation problem of Eq. (8) is given by

$$\begin{bmatrix} \mathbf{p}^*\\ \boldsymbol{\theta}^* \end{bmatrix} = (\mathbf{A}^T \mathbf{R}_{\boldsymbol{\gamma}}^{-1} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{R}_{\boldsymbol{\gamma}}^{-1} \boldsymbol{\gamma}.$$
 (10)

Note that Eq. (10) involves the inversion of a large sparse matrix \mathbf{R}_{γ} which may not be suitable for implementation due to complexity and potential numerical issues. However, this inversion is easily avoided by analytically computing the information matrix $\mathbf{\Omega}_{\gamma} = \mathbf{R}_{\gamma}^{-1}$ using block-matrix inversion rules.

D. Analysis

Small errors in relative orientation estimates add up over time to create rapid growth in position error when unchecked with absolute orientation measurements. This problem arises due to the non-linear nature (trigonometric functions) of orientation. We now proceed to first present an analysis of growth in position error due to noisy relative orientation measurements followed by a discussion of the accuracy of orientation estimation itself.

Let $\mathbf{x}_0 = [0, 0, 0]^T$ be the pose of the robot at time t_0 which is known. Let ${}^{l}\boldsymbol{\Delta}_{ij}$ be the local relative position measurement and $\delta\hat{\theta}_{ij}$ to be the relative orientation measurement between poses x_i and x_j . At each relative orientation measurement, let $\delta\tilde{\theta}_{ij}$ be a small error, hence $\delta\hat{\theta}_{ij} = \delta\theta_{ij} + \delta\tilde{\theta}_{ij}$.

a small error, hence $\delta \hat{\theta}_{ij} = \delta \theta_{ij} + \delta \tilde{\theta}_{ij}$. Then, 2D position \mathbf{p}_n of pose x_n is given by, $\mathbf{p}_n = \mathbf{x}_0 + \sum_{i=0}^{n-1} {}^w \mathbf{\Delta}_{i,i+1}$.

The above equation can be written in terms of the local relative pose transformations as

$$\mathbf{p}_n = \mathbf{x}_0 + \sum_{k=1}^n \left[\left(\prod_{i=0}^{i=k-2} \mathbf{R}(\delta \theta_{i,i+1}) \right)^{l} \mathbf{\Delta}_{k-1,k} \right].$$
(11)

1) Noisy Relative Orientation Measurements: We proceed to analyze the case when only relative orientation measurements are noisy. For clarity of presentation and without loss of generality; (i) we drop the pose subscript and use the fact that x_0 is the origin, (ii) we assume that relative rotation at each step for the true motion of the robot is fixed, thus we fix $\mathbf{R}(\delta\theta_{i,i+1})$ as $\mathbf{R}(\delta\theta)$ and (iii) the relative linear displacement ${}^{l}\boldsymbol{\Delta}_{n-1,n}$ is fixed to $\boldsymbol{\Delta}$ and is known perfectly. Then using Eq. (11) we have $\mathbf{p}_n = \sum_{k=1}^{n} (\mathbf{R}(\delta\theta))^{k-1} \boldsymbol{\Delta}$. Assuming that the only measurement error is in the relative orientation information, the error \mathbf{e}_p at pose x_n is $\mathbf{e}_{p_n} = \sum_{k=1}^{n} ((\mathbf{R}(\delta\theta)\mathbf{R}(\delta\tilde{\theta}))^{k-1} - (\mathbf{R}(\delta\theta))^{k-1})\boldsymbol{\Delta}$. Setting $\mathbf{R}(\delta\theta) = \mathbf{I}$ without loss of generality, we have



Fig. 3. Numerical results for analysis in Section IV-D for 50 Monte Carlo simulations with only relative orientation measurement noise. The figure depicts average terminal position error with $\sqrt{trace(Var(\mathbf{e_p}))}$ bounds for heading noise $\sigma_{\delta\bar{\theta}} = 0.05^{\circ}$. The trajectory length was varied from 1 m to 100 km. Note that at 100 km the $\sqrt{trace(Var(\mathbf{e_p}))}$ bound is \approx 30 km which is 30% of the distance traveled.

 $\mathbf{e}_{p_n} = \sum_{k=1}^n ((\mathbf{R}(\delta \tilde{\theta}))^{k-1} - \mathbf{I}) \boldsymbol{\Delta}$. Assuming small noise, i.e., small $\delta \tilde{\theta}$, and using $cos(\delta \tilde{\theta}) \approx 1$, $sin(\delta \tilde{\theta}) \approx \delta \tilde{\theta}$, we get

$$\mathbf{e}_{p_n} = \sum_{k=1}^n \begin{bmatrix} 0 & -(k-1)\delta\tilde{\theta} \\ (k-1)\delta\tilde{\theta} & 0 \end{bmatrix} \mathbf{\Delta}.$$
 (12)

Setting $\Delta = [1, 1]^T$, and using the fact that sum of first *n* natural numbers is n(n + 1)/2, we get

$$\mathbf{e}_{p_n} = \begin{bmatrix} -\frac{(n-1)(n-2)}{2}\delta\tilde{\theta}\\ \frac{(n-1)(n-2)}{2}\delta\tilde{\theta} \end{bmatrix} \boldsymbol{\Delta}.$$
 (13)

Thus the error in position grows quadratically as the robot moves, Fig. 3 confirms this with empirical results. For $\Delta = [1, 1]^T$, Fig. 3 shows the variation in terminal position error as the trajectory length is increased from 1m to 100km for a small heading error standard deviation of 0.05°. We draw attention to the fact that position error grows super-linearly (quadratic) as the robot moves. Thus, if heading estimates are not fixed with accurate global measurements, the SLAM-front end is bound to drift such that it may not be able to detect loop closure.

2) Noisy Relative Position Measurements: We proceed to analyze the case when only relative position measurements are noisy. Using Eq. (11) and for clarity of presentation and without loss of generality; (i) we drop the pose subscript and use the fact that x_0 is the origin, (ii) we assume that relative rotation at each step for the true motion of the robot is perfectly known thus we fix $\mathbf{R}(\delta\theta_{i,i+1})$ as $\mathbf{R}(\delta\theta) = I$. Then we have $\mathbf{p}_n = \sum_{k=1}^n \Delta_{k-1,k}$. With noisy measurements we have $\hat{\mathbf{p}}_n = \sum_{k=1}^n (\Delta_{k-1,k} + \tilde{\Delta}_{k-1,k})$, where $\tilde{\Delta}_{k-1,k}$ is the zero-mean Gaussian noise in relative position measurements. The position estimation error \mathbf{e}_p at the *n*-th pose is $\mathbf{e}_n = \sum_{k=1}^n (\tilde{\Delta}_{k-1,k})$. Setting $\tilde{\Delta}_{k-1,k} = \tilde{\Delta}$, we can re-write $\mathbf{e}_n = n\tilde{\Delta}$. Thus the error in position grows linearly as the robot moves. We contrast this with the case where we have only relative orientation noise wherein position error grows quadratically.



Fig. 4. Numerical results for analysis in Section IV-D for 50 Monte Carlo simulations with only relative position measurement noise. The figure depicts average terminal position error with $\sqrt{trace(Var(\mathbf{e_p}))}$ bounds for translation noise $\sigma_{\tilde{\Delta}_x} = \sigma_{\tilde{\Delta}_y} = 1$ m. The trajectory length was varied from 1 m to 100 km. Note that at 100 km the $\sqrt{trace(Var(\mathbf{e_p}))}$ bound is \approx 500 m which is 0.5% of the distance traveled.

Figure 4 shows the variation in terminal position error as the trajectory length is increased from 1m to 100 km for a relatively high relative position error (standard deviation of position error is 1m which is equivalent one-step motion of the robot). We draw attention to the fact that position error growth is linear with noisy relative position measurements whereas it is quadratic with noisy relative orientation measurements. Note that the $\sqrt{trace(Var(\mathbf{e_p}))}$ bound is 0.5% of the distance traveled at 100km whereas with small relative orientation noise of 0.05° the same bound is $\approx 30\%$ of the distance moved. This clearly indicates that; (i) the cumulative effect of small deviations in relative orientation measurements results in unreliable localization and (ii) localization uncertainty has the capability to render data-associations for global loop closure detection quite unreliable.

We can see from the simple analysis above that position error grows linearly in the distance moved and there is no (quadratic) compounding of error due to orientation if absolute orientation measurements are available.

V. SIMULATION RESULTS

We tested our approach in multiple high fidelity simulations. To simulate a realistic scenario we obtained the floor plan for a large local HEB grocery store with a 78,000 square feet footprint with a 36 foot high ceiling in College Station, Texas. Using the floor plan we constructed a virtual store environment in 3D and drove a turtlebot equipped with a 2D Lidar, gyro, wheel encoders and a ceiling facing camera (orientation sensing) using Gazebo in ROS. Figure 5(a) shows a virtual turtlebot operation in the simulated environment. To generate the simulation results, we created a series of waypoints through the building which the robot must follow during each Monte Carlo run.

A. Heading Estimation

Figure 5(b) shows pictures taken at the real-world facility. The ceiling has distinctly visible corrugation patterns and



Fig. 5. (a) An instance of a virtual turtlebot operating in the simulated environment. (b) The ceiling of the actual physical grocery store on which the virtual environment was modeled.



Fig. 6. The ceiling direction estimation process for an image captured from the ceiling camera in the simulation study. (a) shows a view of the ceiling in the simulated grocery store model. In (b) we see results of the line detection algorithm (red) and the calculated ceiling direction (green). In (c) the ceiling direction measurements are plotted on polar histogram. (d) shows the zoomed-in histogram at the plot.

structured light fixtures. The ceiling texture of the virtual grocery store replicates the corrugation pattern observed in the physical building. A ceiling facing camera is used to detect line features arising from the corrugated roofing. We use the heading estimation method as described in Section IV-A. Figure 6 shows images of the simulated ceiling and a polar plot of the ceiling direction histogram. We use a bin size $b = 0.1^{\circ}$ for the histogram operation described in Section IV-A. It may be noted here that the lightning in the simulation is uniform, whereas the light sources in the actual store are discretely located from each other.

B. Mapping and Localization

Table I shows the RMS localization error for a total of 200 runs with two different LiDAR ranges of 20m and 30 m. With our approach the average RMS position error was 1 order of magnitude smaller than the standard pose-graph SLAM approach. The results in Table I show that while the standard approach suffered a catastrophic failure $\approx 40\% - 50\%$ of the times **our approach never resulted in a failure**. Fig. 1 shows a comparison of maps created for 1 particular run. As can be seen from these results, the proposed approach constitutes a significant robustness improvement in mapping over the current state of the



Fig. 7. Simulation results for one experiment with 30m range. The figure shows mapping failure as a result of faulty loop closure detection.



Fig. 8. (a) Side-view of the physical robot used in real-world warehouse experiment. (b) Robot's view from its forward facing camera. The ceiling has clearly visible lighting fixtures which are used as orientation information sources.

art. Figure 7 shows mapping failure and localization error for one simulation run with existing state-of-the-art methods. The scan matching algorithm detects an incorrect loop closure due to drift in localization which corrupts the map and the localization diverges rapidly from the ground truth. Such errors are a commonly occurring reason for mapping failure in scan matching based techniques.

VI. PHYSICAL EXPERIMENT RESULTS

We conducted physical experiments with a retrofitted turtlebot to demonstrate the reliability of our mapping algorithm in a warehouse, which is a commonly occurring industrial robotics environment. The facility used was Texas A&M University's Surplus warehouse located at 957 Agronomy Rd, College Station, Texas. The facility has a total floor area of 66,000 sq. ft. The building measures 95 m \times 64 m and is a dynamic environment with constantly moving people, vehicles (vans, forklifts etc.) and goods.

A. Robot Setup

The robot used in our experiments is shown in Fig. 8(a). It is equipped with easily available low-cost commercial sensors including; (i) a 360° LiDAR (RP Lidar) with a range of 5 m, (ii) a monocular ceiling facing camera (Pt Grey Chameleon) is used, the camera runs at 30 Hz, with an image resolution of 640×480 and fields of view being 77.8 and 62.25 degrees along its x and y

(a) Ceiling direction (b) Polar histogram (c) Complete polar mean $\theta_c = 138.09^\circ$, (zoomed-in) histogram $\sigma_c = 0.5479^\circ$.

Fig. 9. View of the ceiling from the upward facing camera mounted on our platform. The orientation sensing system uses straight line light fixtures to detect line features and find the principal direction of the ceiling in the image frame.

axes respectively, (iii) a MEMS IMU, of which only the yaw rate gyro is used and (iv) a commercial off-the-shelf laptop (Core-i7 2.4 GHz, 12 GB RAM). The objective with this robot setup was to show that we are able to achieve highly reliable mapping using low-cost commercial sensors for large scale environments. We drove the robot around the facility and used ROS bags to store the data. The data was then post-processed using our approach and state-of-the-art. Thus identical sensor data was used for both estimates.

B. Heading Estimation

An initial survey of the warehouse facility revealed that the ceiling was equipped with rectangular light fixtures at regular intervals which we decided to leverage for orientation estimation. For the purposes of our experiment, we threshold the ceiling image such that we get a binary image, thus only ceiling lights appear as rectangular bright spots while rest of the image appears black. After thresholding we follow the process described in Algorithm 1. Fig. 9 shows the ceiling camera's view and the thresholded binary image along with heading detection. Heading estimates were computed at 30 Hz. We use a bin size $b = 0.1^{\circ}$ for the histogram operation described in Section IV-A.

1) Orientation Estimation Accuracy: Assuming no inherent bias, heading accuracy is dependent on the number of feature measurements available from the ceiling. More measurements result in better estimation (Law of Large Numbers), this can be seen in the contrast of the polar plots (Fig. 6(c)) for the simulation where we observed ≈ 1000 features in the heaviest bin vs. the real-world experiment (Fig. 9(c)) where we observed 2 features in the heaviest bin. In the simulation we are able to achieve orientation estimation with a standard deviation of $\approx < 0.1^{\circ}$ whereas for the real-world case the number is $\approx 0.5^{\circ}$. As an example, when the simulated ceiling height was lowered to 18 feet, the number of measurements drops to ≈ 100 as opposed to ≈ 1000 available when the ceiling is at a height of 36 feet. This has the direct effect of reducing the accuracy of heading measurements.

C. Mapping

Fig. 10 shows mapping results for the physical experiment. Fig. 10(a) shows the online map as computed by the standard correlative scan matching based SLAM approach. The map suffers a catastrophic failure characterized by a twist in the estimated geometry of the environment. Fig. 10(b) shows the map as estimated by our approach. Our map is globally consistent



Fig. 10. Mapping results for physical warehouse experiment. (a) Incorrectly estimated map with a standard SLAM approach using correlative scan matching in the front-end and G2O in the back-end. (b) Estimated map with our approach.

and does not suffer major twisting whereas the map estimated with the state-of-the-art approach is bent out of shape and unusable for navigation in its current form. A key difference that may be noted is; the edges of environment geometry without using orientation appear sharper (Fig. 10(a)) in contrast to the geometry as estimated when absolute orientation sensing is used (Fig. 10(b)). Heading estimation noise is the driving factor behind this local *map smudging*. Thus higher accuracy in heading estimation should eliminate this issue. We note that in simulations this issue is not observed as heading estimation is more accurate (see Section VI-B1).

VII. CONCLUSIONS

Our method is the first of its kind to fuse absolute orientation sensing in a 2D pose-graph formulation. Through a detailed analysis we showed when relative orientation noise is present, position error grows quadratically in the distance moved. With large localization error bounds one cannot reasonably expect detecting accurate loop closures. Through extensive simulations we have shown that our approach is an order of magnitude more accurate and never fails in testing whereas existing mapping methods fail $\approx 40\% - 50\%$ of the times. Further, we confirmed the robustness of our approach in a physical experiment wherein we mapped a warehouse facility and showed that a regular scan matching approach fails to estimate a consistent map. A key limitation to our current approach is the reliance on structural cues, e.g., straight lines to provide orientation information. In environments which lack distinct heading cues such as airports which can often have large curved atria this approach may not be feasible. An important aspect of future work is to transfer these concepts to various other environment types which may require novel solutions to the orientation estimation problem.

ACKNOWLEDGEMENT

The material presented in this work are subject to pending patent approvals (patent filed). For licensing, please contact Dr. Ismail Sheikh (smismail@tamu.edu) at Texas A&M University Technology Commercialization, 800 Raymond Stotzer Parkway, Suite 2020, College Station, Texas 77845.

REFERENCES

- E. B. Olson, "Real-time correlative scan matching," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2009, pp. 4387–4393.
- [2] Open Karto. [Online]. Available: http://wiki.ros.org/open_karto
- [3] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "G20: A general framework for graph optimization," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2011, pp. 3607–3613.
- [4] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," Auton. Robots, vol. 4, no. 4, pp. 333–349, Oct. 1997.
- [5] F. Dellaert and M. Kaess, "Square root SAM: Simultaneous localization and mapping via square root information smoothing," *Int. J. Robot. Res.*, vol. 25, no. 12, pp. 1181–1203, 2006.
- [6] K. Konolige, G. Grisetti, R. Kümmerle, W. Burgard, B. Limketkai, and R. Vincent, "Efficient sparse pose adjustment for 2D mapping," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Oct. 2010, pp. 22–29.
- [7] J. Folkesson and H. Christensen, "Graphical SLAM—A self-correcting map," in *Proc. IEEE Int. Conf. Robot. Automat.*, Apr. 2004, vol. 1, pp. 383–390.
- [8] I. Mahon, S. B. Williams, O. Pizarro, and M. Johnson-Roberson, "Efficient view-based SLAM using visual loop closures," *IEEE Trans. Robot.*, vol. 24, no. 5, pp. 1002–1014, Oct. 2008.
- [9] J. L. Martínez, J. Morales, A. Mandow, and A. García-Cerezo, "Incremental closed-form solution to globally consistent 2D range scan mapping with two-step pose estimation," in *Proc. 11th IEEE Int. Workshop Adv. Motion Control*, 2010, pp. 252–257.
- [10] L. Carlone, R. Aragues, J. A. Castellanos, and B. Bona, "A fast and accurate approximation for planar pose graph optimization," *Int. J. Robot. Res.*, vol. 33, no. 7, pp. 965–987, 2014.

- [11] L. Carlone and A. Censi, "From angular manifolds to the integer lattice: Guaranteed orientation estimation with application to pose graph optimization," *IEEE Trans. Robot.*, vol. 30, no. 2, pp. 475–492, Apr. 2014.
- [12] K. Khosoussi, S. Huang, and G. Dissanayake, "Exploiting the separable structure of SLAM," in *Proc. Robot., Sci. Syst.*, Rome, Italy, Jul. 13–17, 2015.
- [13] S. Agarwal, V. Shree, and S. Chakravorty, "RFM-SLAM: Exploiting relative feature measurements to separate orientation and position estimation in SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2017, pp. 6307–6314, doi: 10.1109/ICRA.2017.7989746.
- [14] N. Boumal, A. Singer, and P. A. Absil, "Robust estimation of rotations from relative measurements by maximum likelihood," in *Proc. 52nd IEEE Conf. Decis. Control*, Dec. 2013, pp. 1156–1161.
- [15] L. Carlone, R. Tron, K. Daniilidis, and F. Dellaert, "Initialization techniques for 3D SLAM: A survey on rotation estimation and its use in pose graph optimization," in *Proc. IEEE Int. Conf. Robot. Automat.*, May 2015, pp. 4597–4604.
- [16] H. Wang, W. Mou, G. Seet, M.-H. Li, M. W. S. Lau, and D.-W. Wang, "Real-time visual odometry estimation based on principal direction detection on ceiling vision," *Int. J. Automat. Comput.*, vol. 10, no. 5, pp. 397–404, Oct. 2013. [Online]. Available: https://doi.org/10.1007/s11633-013-0736-7
- [17] T. Duckett, S. Marsland, and J. Shapiro, "Fast online learning of globally consistent maps," *Auton. Robots*, vol. 12, no. 6, pp. 287–300, May 2002.