On the Treatment of Relative-Pose Measurements for Mobile Robot Localization

Anastasios I. Mourikis and Stergios I. Roumeliotis
Dept. of Computer Science & Engineering, University of Minnesota, Minneapolis, MN 55455
Email: {mourikis|stergios}@cs.umn.edu

***Abstract*— In this paper, we study the problem of localization using relative-state estimates. It is shown, that when the same exteroceptive sensor measurement is processed for the computation of two consecutive displacement estimates (both forward and backward in time) these estimates are correlated, and an analysis of the exact structure of the correlations is performed. This analysis is utilized in the design of data fusion algorithms, that correctly account for the existing correlations. We examine two cases: i) pose propagation based exclusively on inferred displacement measurements, and ii) the fusion of proprioceptive sensor information with relative-state measurements. For the latter case, an efficient EKF-based estimation algorithm is proposed, that extends the approach of [1]. Extensive simulation and experimental results are presented, that verify the validity of the presented method.**

I. Introduction

Accurate localization is a prerequisite for a robot to be able to interact with its environment in a meaningful way. The most commonly available sensors for acquiring localization information are proprioceptive sensors, such as wheel encoders, gyroscopes, and accelerometers, that provide information about the robot's motion. By integrating proprioceptive measurements over time, it is possible to estimate the total displacement from a starting point, and this method of localization is often called Dead Reckoning (DR) [2]. The limitation of DR is that since no external reference signals are employed for correction, estimation errors accumulate over time, and the pose estimates drift from their real values. In order to improve the accuracy of localization, most algorithms fuse the proprioceptive information with data from exteroceptive sensors (such as cameras [3], [4], laser range finders [5] sonars [6], etc) that provide measurements of parameters of interest in the environment.

When an exteroceptive sensor measures the position of a set of features with respect to the robot at two different time instants, then it is possible (under necessary observability assumptions) to create an *inferred* measurement of the robot's displacement. Examples of algorithms that process exteroceptive data to infer motion include laser scan matching [5], [7], vision-based motion estimation techniques using stereoscopic [3], [4], and monocular [8] image sequences, and matching of sonar returns [6]. The inferred *relative-state* measurements that are produced can either be integrated over time to provide pose estimates for the robot's state at each time instant [4], or fused with proprioceptive sensory input, in order to benefit from both available sources of positioning information [1].

A characteristic which is common in most cases where exteroceptive measurements are employed to infer robot displacement is that consecutive relative-state measurements are stochastically *correlated*. The correlation is introduced from the fact that the measurements recorded at time-step k (e.g., the relative positions of landmarks) are used in order to estimate the displacement during the time intervals $[k-1, k]$ *and* $[k, k+1]$. As a result, any errors in the measurements at time step k affect both displacement estimates, thereby rendering them correlated. If the measurements are treated as being uncorrelated (as is customarily done [9], [1], [3]) information is lost, and the estimates for the robot's state and covariance are not optimal. This fact has been generally overlooked in the literature, and, to the best of our knowledge, no prior work exists that directly addresses this issue.

In this paper, we propose a direct approach to the processing of correlated displacement measurements, that extends the *Stochastic Cloning* Kalman Filter (SC-KF) algorithm of [1]. In particular, in Section III we show how the correlation of the measurements can be accounted for, when propagating the covariance of the robot's pose in time, based only on relative-state measurements. Additionally, in Section IV we propose a formulation of the Extended Kalman Filter (EKF) for fusing proprioceptive and relative-state measurements, that correctly treats correlation between consecutive measurements. This is achieved by *augmenting* the state vector to include the *measurement errors* for each exteroceptive measurement. The simulation and experimental results from the application of this method demonstrate that correctly accounting for the correlations results in better state estimates, as well as in covariance estimates that reflect the true uncertainty in the robot's pose more accurately.

II. Related Work

Once a displacement measurement is derived from an exteroceptive sensor, in most cases it must be combined with other position estimates derived from onboard proprioceptive sensors. An appealing solution to this problem is to use the previous position estimates for converting the relative-pose measurements to absolute position *pseudo-measurements*, and treat them as such [10]. However, this approach is correct only when the orientation of the vehicle is precisely known. Moreover, to guarantee consistency of the pose estimates, the covariance matrix of the pseudo-measurements has to be artificially inflated [11].

A difficulty that arises when processing relative-pose measurements is the existence of correlations between consecutive measurement errors (cf. Section III). A simplistic approach to this problem would be to discard correlations, by separating the measurements recorded at each robot pose in two non-overlapping sets; one is used to estimate motion forward in time, while the second is used to estimate displacement backward in time. This solution, however, would be far from optimal, as it would result in less accurate displacement estimates. An indirect solution is to avoid the introduction of correlations altogether, by *not* using the exteroceptive measurements to infer displacement directly. In this formulation, the robot's pose and the position of the environmental features are estimated *jointly*, thus introducing the well-known Simultaneous Localization and Mapping (SLAM) problem, which has been extensively studied in the robotics community (e.g., [12], [13], [14]). If an exact solution to SLAM was possible, the resulting pose estimates would be optimal, since all the positioning information is used, and all the inter-dependencies between the robot and the feature states are accounted for. However, the major limitation of SLAM is that its computational complexity and memory requirements increase quadratically with the number of features in the environment. This implies that, if a robot operates over extended periods of time, the amount of resources that need to be allocated for localization tend to become unacceptably large, if real-time performance is necessary.

In this paper, we propose an algorithm for fusing proprioceptive information with relative-state measurements, which extends the SC-KF algorithm [1], to be applicable to the case of *correlated* relative-state measurements. In our algorithm, the exteroceptive measurements are considered in pairs of consecutive measurements, that are first processed in order to create an inferred relative-state measurement, and then fused with the proprioceptive measurements. The sole objective of the algorithm is the estimation of the robot's state, and therefore we do *not* estimate the states of features in the environment. Our motivation for this arises from the fact that in applications where building a map is not necessary, the overhead of performing SLAM may not be justified. In cases where real-time performance is required (e.g., autonomous aircraft landing), the proposed algorithm is able to optimally fuse the, potentially correlated, relative-state measurements, with the minimum computational overhead. Before presenting this algorithm, in the following section we analyze the structure of the correlations that exist between consecutive displacement measurements, and demonstrate how these should be treated in the propagation of the pose estimates' uncertainty.

III. Covariance Propagation Based on Displacement Measurements

In this section, we consider the case in which the pose estimate of a robot[1], $\widehat{X}_k$, is propagated in time using only displacement estimates, that are acquired by processing exteroceptive measurements. Let z_k and z_{k+1} denote the vectors of exteroceptive measurements at time-steps k and $k+1$, respectively, whose covariance matrices are denoted as R_k and R_{k+1}. These are, for example, the measurements of the position of landmarks with respect to the robot, or the range measurements of a laser range finder. By processing these measurements (for example, by performing laser scan matching), an estimate, $z_{k/k+1}$, for the change in the robot pose between time-steps k and $k+1$ is computed, which is described by a function (either closed-form or implicit):

$$z_{k/k+1} \quad = \quad \xi_{k/k+1}(z_k, z_{k+1}) \tag{1}$$

Linearization of this last expression enables us to relate the error in the displacement estimate, $\widetilde{z}_{k/k+1}$, to the errors in the exteroceptive measurements:

$$\widetilde{z}_{k/k+1} \quad \simeq \quad J^k_{k/k+1}\widetilde{z}_k + J^{k+1}_{k/k+1}\widetilde{z}_{k+1} \tag{2}$$

Here, we assume that the errors in the exteroceptive measurements, $\widetilde{z}_k$ and $\widetilde{z}_{k+1}$, are zero-mean and independent, an assumption which holds in most practical cases, when proper sensor characterization is performed. In Eq. (2) $J^k_{k/k+1}$ and $J^{k+1}_{k/k+1}$ are the Jacobians of the function $\xi_{k/k+1}(z_k, z_{k+1})$ with respect to z_k and z_{k+1}, respectively, i.e.,

$$J^k_{k/k+1} = \nabla_{z_k}\xi_{k/k+1} \quad \text{and} \quad J^{k+1}_{k/k+1} = \nabla_{z_{k+1}}\xi_{k/k+1}$$

Once the displacement estimate $z_{k/k+1}$ between time-steps k and $k+1$ has been computed, the pose estimate for the robot at time step $k+1$ is evaluated by combining the previous pose estimate and the displacement information, by an appropriate, generally nonlinear function:

$$\widehat{X}_{k+1} \quad = \quad g(\widehat{X}_k, z_{k/k+1}) \tag{3}$$

By linearizing this equation, the pose errors at time step $k+1$ can be related to the error in the previous state estimate and that in the displacement measurement:

$$\widetilde{X}_{k+1} \quad \simeq \quad \Phi_k\widetilde{X}_k + \Gamma_k\widetilde{z}_{k/k+1} \tag{4}$$

where Φ_k and Γ_k represent the Jacobians of the state propagation function, $g(\widehat{X}_k, z_{k/k+1})$, with respect to the previous pose, and the relative pose measurement, respectively. During localization, it is necessary to provide a measure of the quality of the pose estimates, since the uncertainty in the robot's pose should be taken into consideration for motion planning and navigation. For a Gaussian distribution, a sufficient indicator of uncertainty is the covariance matrix of the pose errors:

$$\begin{aligned} P_{k+1} &= E\{\widetilde{X}_{k+1}\widetilde{X}^T_{k+1}\} \\ &= \Phi_k P_k \Phi^T_k + \Gamma_k R_{k/k+1}\Gamma^T_k \\ &+ \Phi_k E\{\widetilde{X}_k\widetilde{z}^T_{k/k+1}\}\Gamma^T_k + \Gamma_k E\{\widetilde{z}_{k/k+1}\widetilde{X}^T_k\}\Phi^T_k \end{aligned} \tag{5}$$

where $R_{k/k+1}$ denotes the covariance matrix of the displacement estimates. If the measurement noise, $\widetilde{z}_{k/k+1}$, and state error, $\widetilde{X}_k$, are uncorrelated, the last two terms in Eq. (5) are equal to zero matrices, and the covariance propagation expression becomes identical to the well-known covariance propagation expression of the Extended Kalman Filter [15]. Although this is a common assumption in the literature (e.g., [1], [9]), we now demonstrate that it does *not* hold in general, and the last two terms in Eq. (5) are *not* equal to zero. In particular, by

[1]Throughout this paper, the "hat" symbol, $\widehat{\ }$, is used to denote the estimated value of a quantity, while the "tilde" symbol, $\widetilde{\ }$, is used to signify the error between the actual value of a quantity and its estimate, i.e., $\widetilde{x} = x - \widehat{x}$.

linearizing the state propagation equation at time-step k, we obtain (cf. Eq. (4)):

$$E\{\widetilde{z}_{k/k+1}\widetilde{X}_k^T\} = E\left\{\widetilde{z}_{k/k+1}\left(\Phi_{k-1}\widetilde{X}_{k-1} + \Gamma_{k-1}\widetilde{z}_{k-1/k}\right)^T\right\}$$
$$= E\{\widetilde{z}_{k/k+1}\widetilde{X}_{k-1}^T\}\Phi_{k-1}^T + E\{\widetilde{z}_{k/k+1}\widetilde{z}_{k-1/k}^T\}\Gamma_{k-1}^T$$
$$= E\{\widetilde{z}_{k/k+1}\widetilde{z}_{k-1/k}^T\}\Gamma_{k-1}^T \quad (6)$$

At this point we note that the error term $\widetilde{X}_{k-1}$ depends on the measurement errors of all exteroceptive measurements up to, and including, time-step $k-1$, while the error term $\widetilde{z}_{k/k+1}$ depends on the measurement errors at time-steps k and $k+1$ (cf. Eq. (2)). As a result, the errors $\widetilde{X}_{k-1}$ and $\widetilde{z}_{k/k+1}$ are independent, and therefore, by applying the zero-mean assumption for the error $\widetilde{z}_{k/k+1}$ we obtain $E\{\widetilde{z}_{k/k+1}\widetilde{X}_{k-1}^T\} = 0$. In order to evaluate the term $E\{\widetilde{z}_{k/k+1}\widetilde{z}_{k-1/k}^T\}$, which expresses the correlation between the consecutive displacement estimates, we employ Eq. (2), and the independence of exteroceptive measurement errors at different time-steps, to obtain:

$$E\{\widetilde{z}_{k/k+1}\widetilde{z}_{k-1/k}^T\} = J_{k/k+1}^k E\{\widetilde{z}_k\widetilde{z}_k^T\} J_{k-1/k}^{k\ T}$$
$$= J_{k/k+1}^k R_k J_{k-1/k}^{k\ T} \quad (7)$$

This result implies that consecutive displacement estimates are *not* independent. However, the statistical correlation between them *is* computable in closed form, and can be accounted for in the propagation of the state covariance matrix. Substituting from Eqs. (6) and (7) into Eq. (5), we obtain the final expression for propagating the pose covariance based on inferred displacement measurements:

$$P_{k+1} = \Phi_k P_k \Phi_k^T + \Gamma_k R_{k/k+1}\Gamma_k^T + D_{k+1} + D_{k+1}^T \quad (8)$$

with

$$D_{k+1} = \Phi_k\Gamma_{k-1}J_{k-1/k}^k R_k J_{k/k+1}^{k\ T}\Gamma_k^T$$

The experimental results presented in Section V-B demonstrate, that by employing this expression for propagating the robot's pose covariance matrix, we are able to compute covariance estimates that accurately represent the robot's uncertainty.

IV. Filtering with correlated relative-state measurements

In this section, we consider the situation in which relative-state measurements are fused with proprioceptive sensory information to estimate the robot's pose. Since the proprioceptive and exteroceptive measurements are received from two independent sources of information, fusing them will always result in superior estimation accuracy, compared to the accuracy attainable when the robot's pose is propagated based solely on one of the two types of measurements.

Two challenges arise when fusing relative-state and proprioceptive measurements: firstly, since each displacement measurement relates the robot's state at two *different* time instants, the "standard" formulation of the EKF, in which the filter's state comprises only the current state of the robot, is not adequate. Secondly, as shown in the preceding section (cf. Eq. (7)), the consecutive displacement measurements are correlated, and this violates one of the basic assumptions of the EKF, that of the independence of the measurements [15]. To address the first challenge, we adopt the approach proposed in [1], that requires the *augmentation* of the EKF (error) state vector[2] to include two copies of the robot's error state (cloning). The first copy represents the pose error at the time instant when the latest exteroceptive measurement was recorded, while the second copy represents the error in the robot's current state. Consequently, the robot states that are related by each displacement estimate are represented explicitly in the filter state.

The second challenge is addressed by further augmenting the state vector to include the *errors* of the latest exteroceptive measurement. Thus, if the most recent exteroceptive measurement was recorded at time-step k, the filter's state vector at time-step $k+i$ is given by[3]

$$\breve{X}_{k+i|k} = \begin{bmatrix}\widetilde{X}_{k|k}^T & \widetilde{X}_{k+i|k}^T & \widetilde{z}_k^T\end{bmatrix}^T \quad (9)$$

By including the measurement error in the state vector of the system, the dependency of the relative-state measurements $z_{k/k+1}$ on the exteroceptive measurements z_k is transformed into a dependency on the *current state of the filter*, and the problem can now be treated in the standard EKF framework. It should be noted, that since the error in the intermediate measurement is the source of the correlation between the current and previous displacement estimates (cf. Eq. (7)), this is the "minimum length" vector that we need to append to the state vector, in order to sufficiently describe the existing dependencies. In the following sections, we present in detail the propagation and update phases of the filter.

A. State propagation

Consider the case where the filter's state covariance matrix, immediately after the exteroceptive measurement z_k has been processed, is given by:

$$\breve{P}_{k|k} = \begin{bmatrix} P_{k|k} & P_{k|k} & P_{X_k z_k} \\ P_{k|k} & P_{k|k} & P_{X_k z_k} \\ P_{X_k z_k}^T & P_{X_k z_k}^T & R_k \end{bmatrix} \quad (10)$$

where $P_{k|k}$ is the covariance of the actual robot pose at time-step k, R_k is the covariance matrix of the error $\widetilde{z}_k$, and $P_{X_k z_k} = E\{\widetilde{X}_k\widetilde{z}_k^T\}$ is the cross-correlation between the robot's state and the measurement error at time-step k (the derivation of a closed-form expression for $P_{X_k z_k}$ is presented in Section IV-B). We note that cloning the robot's state creates two random variables that convey the same information, and hence are fully correlated [1]. This explains the structure of the covariance matrix in Eq. (10).

Between two consecutive updates, the proprioceptive measurements are employed to propagate the filter's state and its covariance. Let the proprioceptive measurement at time-step k be denoted as v_k, and its noise covariance matrix as Q_k.

[2] Since the EKF is employed for estimation, the state vector comprises of the *errors* in the estimated quantities, rather than the estimates. Therefore, cloning has to be applied to both the error states, and the actual estimates.

[3] In the remainder of the paper the subscript $\ell|j$ denotes the estimated value of a quantity at time step ℓ, after exteroceptive measurements up to time-step j, and proprioceptive measurements up to time-step $\ell-1$, have been processed.

The estimate for the robot's pose is propagated in time by the, generally non-linear, equation:

$$\widehat{X}_{k+1|k} = f(\widehat{X}_{k|k}, v_k) \tag{11}$$

Linearization of the last expression yields the error propagation equation for the (evolving) robot state:

$$\widetilde{X}_{k+1|k} \simeq F_k \widetilde{X}_{k|k} + G_k \widetilde{v}_k \tag{12}$$

where F_k and G_k are the Jacobians of $f(\widehat{X}_{k+1|k}, v_k)$ with respect to $\widehat{X}_{k|k}$ and v_k, respectively. Since the cloned state, as well as the estimates for the measurement error $\widetilde{z}_k$ do *not* change with the incorporation of a new proprioceptive measurement, the error propagation equation for the entire state vector is given by

$$\breve{X}_{k+1|k} = \breve{F}_k \breve{X}_{k|k} + \breve{G}_k \widetilde{v}_k \tag{13}$$

$$\text{with} \quad \breve{F}_k = \begin{bmatrix} I & 0 & 0 \\ 0 & F_k & 0 \\ 0 & 0 & I \end{bmatrix} \quad \text{and} \quad \breve{G}_k = \begin{bmatrix} 0 \\ G_k \\ 0 \end{bmatrix} \tag{14}$$

thus the covariance matrix of the propagated filter state is:

$$\breve{P}_{k+1|k} = \breve{F}_k \breve{P}_{k|k} \breve{F}_k^T + \breve{G}_k Q_k \breve{G}_k^T \tag{15}$$

It is straightforward to show by induction that if m propagation steps take place between two consecutive relative-state updates the covariance matrix $\breve{P}_{k+m|k}$ is determined as

$$\breve{P}_{k+m|k} = \begin{bmatrix} P_{k|k} & \mathcal{F}_{k/k+m}^T P_{k|k} & P_{X_k z_k} \\ \mathcal{F}_{k/k+m} P_{k|k} & P_{k+m|k} & \mathcal{F}_{k/k+m} P_{X_k z_k} \\ P_{X_k z_k}^T & P_{X_k z_k}^T \mathcal{F}_{k/k+m}^T & R_k \end{bmatrix} \tag{16}$$

where $\mathcal{F}_{k/k+m} = \prod_{i=0}^{m-1} F_{k+i}$, and $P_{k+m|k}$ is the propagated covariance of the robot state at time-step $k+m$. The last expression indicates that exploiting the structure of the propagation equations allows for the covariance matrix of the filter to be propagated with minimal computation. In an implementation where efficiency is of utmost importance, the product $\mathcal{F}_{k/k+m}$ can be accumulated, and the matrix multiplications necessary to compute the $\breve{P}_{k+m|k}$ can be delayed, and carried out only when a new exteroceptive measurement has to be processed.

B. State update

We now assume that a new exteroceptive measurement, z_{k+m}, is recorded at time-step $k+m$, and is processed along with z_k to produce a relative-state measurement, $z_{k/k+m} = \xi_{k/k+m}(z_k, z_{k+m})$, relating the robot poses X_k and X_{k+m}. It should be pointed out that $z_{k/k+m}$ is *not* required to provide information about all the degrees of freedom of the pose change between times k and $k+m$. This allows for processing relative-state measurements in cases where the complete displacement cannot be determined (e.g., when estimating pose change based on point-feature correspondences with a single camera, the scale is unobservable [8]). Thus, the relative-state measurement is equal to a general function of the robot poses at time-steps k and $k+m$, with the addition of error:

$$z_{k/k+m} = h(X_k, X_{k+m}) + \widetilde{z}_{k/k+m} \tag{17}$$

The expected value of $z_{k/k+m}$ is computed based on the estimates for the state at times k and $k+m$, as $\widehat{z}_{k/k+m} = h(\widehat{X}_{k|k}, \widehat{X}_{k+m|k})$, and therefore the innovation is given by:

$$\begin{aligned} r &= z_{k/k+m} - \widehat{z}_{k/k+m} \\ &\simeq \begin{bmatrix} H_k & H_{k+m} & J_{k/k+m}^k \end{bmatrix} \begin{bmatrix} \widetilde{X}_{k|k} \\ \widetilde{X}_{k+m|k} \\ \widetilde{z}_k \end{bmatrix} + J_{k/k+m}^{k+m} \widetilde{z}_{k+m} \\ &= \breve{H}_{k+m} \breve{X}_{k+m/k} + J_{k/k+m}^{k+m} \widetilde{z}_{k+m} \end{aligned} \tag{18}$$

where H_k (H_{k+m}) is the Jacobian of $h(X_k, X_{k+m})$ with respect to X_k (X_{k+m}), and $J_{k/k+m}^k$ ($J_{k/k+m}^{k+m}$) is the Jacobian of $z_{k/k+m} = \xi_{k/k+m}(z_k, z_{k+m})$ with respect to z_k (z_{k+m}).

The result of Eq. (18) demonstrates that by incorporating the measurement errors, $\widetilde{z}_k$, in the state vector, the only component of the innovation that is not dependent on the state is the measurement noise, $\widetilde{z}_{k+m}$, which is *independent* of all other error terms in Eq. (18). Thus, the Kalman filter equations can be applied to update the state. The covariance of the residual is equal to

$$\breve{S} = \breve{H}_{k+m} \breve{P}_{k+m|k} \breve{H}_{k+m}^T + J_{k/k+m}^{k+m} R_{k+m} J_{k/k+m}^{k+m\,T} \tag{19}$$

while the Kalman gain is computed as:

$$\breve{K} = \breve{P}_{k+m|k} \breve{H}_{k+m}^T \breve{S}^{-1} = \begin{bmatrix} K_k^T & K_{k+m}^T & K_{z_k}^T \end{bmatrix}^T \tag{20}$$

We note that although the measurement z_{k+m} can be used to update the estimates for the robot's pose at time step k and for the measurement error $\widetilde{z}_k$, our goal is to update only the current state of the robot (i.e., the state at time step $k+m$) and its covariance. Therefore, only the corresponding block element K_{k+m} of the Kalman gain matrix needs to be computed. The equation for updating the current robot state is:

$$\widehat{X}_{k+m|k+m} = \widehat{X}_{k+m|k} + K_{k+m} r \tag{21}$$

While the covariance matrix of the updated robot state is:

$$P_{k+m|k+m} = P_{k+m|k} - K_{k+m} \breve{S} K_{k+m}^T \tag{22}$$

The final step in the process of updating the filter state is to evaluate the new augmented covariance matrix, that will be required for processing the next relative-state measurement. Immediately after $z_{k/k+m}$ is processed, the clone of the previous state error, $\widetilde{X}_{k|k}$, and the previous measurement error, $\widetilde{z}_k$, are discarded. The robot's state at the current time-step, $X_{k+m|k+m}$, is cloned, and the exteroceptive measurement errors, $\widetilde{z}_{k+m}$, are appended to the new filter state. Thus, the filter error-state vector becomes

$$\breve{X}_{k+m|k+m} = \begin{bmatrix} \widetilde{X}_{k+m|k+m}^T & \widetilde{X}_{k+m|k+m}^T & \widetilde{z}_{k+m}^T \end{bmatrix}^T \tag{23}$$

To compute the new filter covariance matrix $\breve{P}_{k+m|k+m}$, the correlation between the robot's error state, $\widetilde{X}_{k+m|k+m}$, and the measurement error vector, $\widetilde{z}_{k+m}$, has to be determined. From Eq. (21) we obtain:

$$\widetilde{X}_{k+m|k+m} = \widetilde{X}_{k+m|k} - K_{k+m} r \tag{24}$$

and employing the result of Eq. (18) yields:

$$\begin{aligned} P_{X_{k+m}z_{k+m}} &= E\{\left(\widetilde{X}_{k+m|k} - K_{k+m}r\right)\widetilde{z}_{k+m}^T\} \\ &= -K_{k+m}E\{r\widetilde{z}_{k+m}^T\} \\ &= -K_{k+m}J_{k/k+m}^{k+m}R_{k+m} \end{aligned} \quad (25)$$

In this derivation, the statistical independence of the error $\widetilde{z}_{k+m}$ to the errors in the state $\breve{X}_{k+m|k}$ has been employed. Using this result, the covariance matrix of the augmented state at time $k+m$ has the same structure as the matrix in Eq. (10) (for indices $k+m$ instead of k).

C. Discussion

From the preceding presentation it becomes apparent that the augmentation of the covariance matrix, that is employed in order to correctly treat the correlations between the consecutive relative-state measurements, inflicts an overhead in terms of computation and memory requirements, which may become cumbersome if the dimension of the measurement vector at time-step k, M_k, is larger than the dimension of the robot's state, N. If the correlations are ignored, as in [1], the size of the state vector in the filter equals double the size of the robot's state, and the computational complexity, as well as the memory requirements of the filter are $\mathcal{O}(N^2)$. In the algorithm proposed in this paper, the most computationally expensive operation, for $M_k \gg N$, is the evaluation of the covariance matrix of the residual (Eq. (19)). Since $\breve{P}_{k+m|k}$ is of dimension $2N+M_k$, the computational complexity of obtaining $\breve{S}$ is generally $\mathcal{O}((2N+M_k)^2) \approx \mathcal{O}(N^2+M_k^2)$. However, in most cases, the vector of exteroceptive measurements commonly comprises a relatively small number of features, detected in the robot's vicinity, e.g., the relative positions of landmarks, the image coordinates of visual features, or the range measurements at specific angles. In such cases, the measurements of the individual features are mutually *independent*, and therefore the covariance matrices R_k and R_{k+m} are block diagonal. By exploiting the structure of $\breve{P}_{k+m|k}$ in this situation, the computational complexity of evaluating Eq. (19) becomes $\mathcal{O}(N^2+M_k)$. Moreover, when the matrices R_k and R_{k+m} are block diagonal, the covariance matrix $\breve{P}_{k+m|k}$ is sparse, which reduces the storage requirements of the algorithm to $\mathcal{O}(N^2+M_k)$.

These complexity figures should be compared to the complexity of performing SLAM, which, as discussed in Section II, is an alternative solution to the problem of processing correlated relative-state measurements. The complexity of performing SLAM in the classic EKF formulation is quadratic in the total number of features included in the state vector. In most cases this number is orders of magnitude larger compared to the number of features detected in each location. Even if an approximate SLAM algorithm is used (e.g., [13], [14]), the largest proportion of the robot's computational resources are devoted to maintaining a constantly enlarging map. This may not be necessary, when only the robot's pose estimates are of interest for a given application.

Additionally, SLAM requires that the states of the features be completely observable, in order for these to be included in the state vector. In cases where a single measurement does not provide sufficient information to initialize a feature's position estimate with bounded uncertainty, complicated feature initialization schemes need to be implemented [16], [17]. In contrast, in the proposed method feature initialization is not required, since the *measurement errors*, which are not explicitly estimated, are included in the augmented state vector.

Furthermore, since in the SC-KF formulation, only *pairs* of consecutive sets of exteroceptive measurements are considered, the data association problem is simplified. In SLAM, correspondence search has to be performed with *all* map features in the robot's vicinity. Thus, the computational overhead is considerably higher [18]. To facilitate robust data association, it is common practice to employ a feature detection algorithm that processes the raw sensor data to extract "high-level" features (e.g., landmarks such as corners, junctions, straight-line segments, distinctive image features). Then, *only* these features are employed for SLAM.

Extracting high-level features results in more robust and computationally tractable algorithms (e.g., laser scans consist of hundreds of points, but only a few corner features are usually present in each scan). This approach, however, effectively *discards information* contained in the "low-level" sensor data. Consequently, the resulting estimates for the robot's pose are suboptimal, compared to the estimates that would be obtained if all available information was utilized. Maintaining and processing the entire history of raw sensor input (e.g., [19]) can clearly lead to excellent localization performance, but with the currently available computing capabilities of robots, this task cannot be performed in real-time. One advantage of the SC-KF approach is that it can utilize all information that exists in two *consecutive* exteroceptive measurements (i.e., most laser points in two scans can be used to estimate displacement by laser scan matching).

At this point, it should be made clear that the positioning accuracy obtained when only *pairs* of exteroceptive measurements are considered is inferior to that of SLAM, as *no loop closing* occurs. Essentially, the SC-KF approach offers an "enhanced" form of Dead Reckoning, in the sense that the uncertainty of the robot's state monotonically increases over time. The rate of increase, though, is significantly lower compared to that attained when only proprioceptive measurements are used. However, we note that in the SC-KF approach the state vector X_k is *not* required to contain only the robot's pose. If high-level, stable features (landmarks) are available, that can be used for SLAM, their positions can be included in the state vector X_k. Therefore, the SC-KF method for processing relative-state measurements can be expanded and integrated with the SLAM framework. This would further improve the attainable localization accuracy within areas with lengthy loops. Since this modification is beyond the scope of this work, in the following section we present experimental results applying the SC-KF methodology for the case where only relative-state and proprioceptive measurements are considered.

V. Experimental results

For the experiments, a Pioneer II robot equipped with a laser rangefinder has been used. The robot's pose comprises

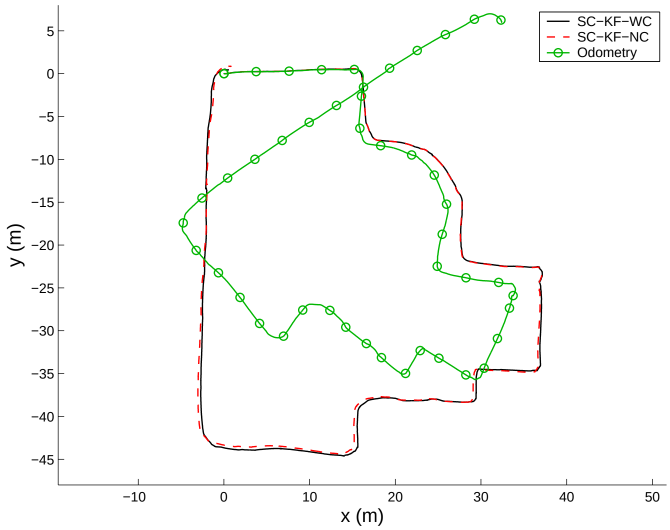

Fig. 1. The estimated trajectory of the robot using the SC-KF-WC algorithm (solid line), the SC-KF-NC algorithm (dashed line), and odometry only (solid line with circles).

its position and orientation in the global frame:

$$X_k = \begin{bmatrix} {}^G x_k & {}^G y_k & {}^G \phi_k \end{bmatrix}^T = \begin{bmatrix} {}^G p_k^T & {}^G \phi_k \end{bmatrix}^T \quad (26)$$

We first present results from the application of the SC-KF and then study the case where the robot's state is propagated based on displacement estimates exclusively (i.e., no proprioceptive measurements are processed).

A. Stochastic Cloning Kalman Filter

In this experiment, odometry measurements are fused with displacement estimates that are obtained by laser scan matching with the method presented in [7]. The robot traversed a trajectory of approximately 165m, while recording 378 laser scans. We here compare the performance of the SC-KF algorithm presented in this paper, that correctly accounts for temporal correlations in the displacement measurements, to that of [1], where correlations are ignored. The two algorithms are referred to as SC-KF-WC (with correlations) and SC-KF-NC (no correlations), respectively.

The estimated robot trajectories resulting from the application of the two algorithms, as well as the trajectory based on odometry only, are shown in Fig. 1. Additionally, in Fig. 2, we present the time evolution of the covariance estimates for the robot pose. We observe that correctly accounting for the correlations between consecutive displacement estimates in the SC-KF, results in smaller covariance values. Even though ground truth for the entire trajectory is not known, the final robot pose is known to coincide with the initial one. The errors in the final robot pose are equal to $\widetilde{X} = [0.5m \;\; 0.44m \;\; -0.11^o]^T$ (0.4% of the trajectory length) for the SC-KF-WC, $\widetilde{X} = [0.61m \;\; 0.65m \;\; -0.13^o]^T$ (0.54% of the trajectory length) for the SC-KF-NC, and $\widetilde{X} = [32.4m \;\; 5.95m \;\; -69.9^o]^T$ (19.9% of the trajectory length) for Dead Reckoning based on odometry. From these error values, as well as from visual inspection of the trajectory estimates in Fig. 1, we conclude that both the SC-KF-WC and the SC-KF-NC yield very similar results.

1) Impact of correlations: Clearly, the lack of ground truth data along the entire trajectory for the real-world experiment does not allow for a detailed comparison of the performance of the SC-KF-WC and SC-KF-NC algorithms; both appear to attain comparable estimation accuracy. In order to perform a more thorough assessment of the impact of the measurement correlations on the position accuracy and the uncertainty estimates, simulation experiments have also been conducted. The primary objective of these experiments is to study the behavior of the estimation errors as compared to the computed covariance values, when the correlations between consecutive measurements are accounted for, vs. when they are ignored.

For the simulation results shown here, a robot moves in a circular trajectory of radius 4m, while observing a wall that lies 6m from the center of its trajectory. The relative-pose measurements in this case are created by performing line-matching, instead of point matching between consecutive scans [20]. Since only one line is available, the motion of the robot along the line direction is unobservable. To avoid numerical instability in the filter, the displacement measurements $z_{k/k+m}$, computed by line-matching are projected onto the observable subspace, thus creating a relative-state measurement of dimension 2.

In Fig. 3, the robot pose errors (solid lines) are shown, along with the corresponding $99,8\%$ percentile of their distribution (dashed lines with circles). The left column shows the results for the SC-KF-WC algorithm presented in Section IV, while the right one for the SC-KF-NC algorithm. As evident from Fig. 3, the covariance estimates of the SC-KF-NC are not commensurate with the corresponding errors. When the temporal correlations of the measurements are properly treated, as is the case for the SC-KF-WC, substantially more accurate covariance estimates, that reflect the true uncertainty of the robot's state, are computed. Moreover, evaluation of the rms value of the pose errors shows that the errors for the SC-KF-WC algorithm, which accounts for the correlations, are 25% smaller compared to those of the SC-KF-NC.

B. State Propagation based on Displacement Estimates

In this Section, we present results for the case in which the robot's pose is estimated using *only* displacement estimates computed from laser scan matching. In Fig. 4, we plot the estimated robot trajectory, along with the map of the area, constructed by overlaying all the scan points, transformed using the estimates of the robot pose (we stress that the map is only plotted for visualization purposes, and is *not* estimated by the algorithm). For this experiment we used the same dataset used for the experiments in the previous section. In Fig. 5, the covariance estimates for the robot's pose, computed using Eq. (8), are presented (solid lines) and compared to those computed when the correlations between the consecutive displacement estimates are ignored (dashed lines). As expected, the pose covariance is larger when only displacement measurements are used, compared to the case where odometry measurements are fused with displacement measurements (cf. Fig. 2). From Fig. 5 we observe that accounting for the correlations results

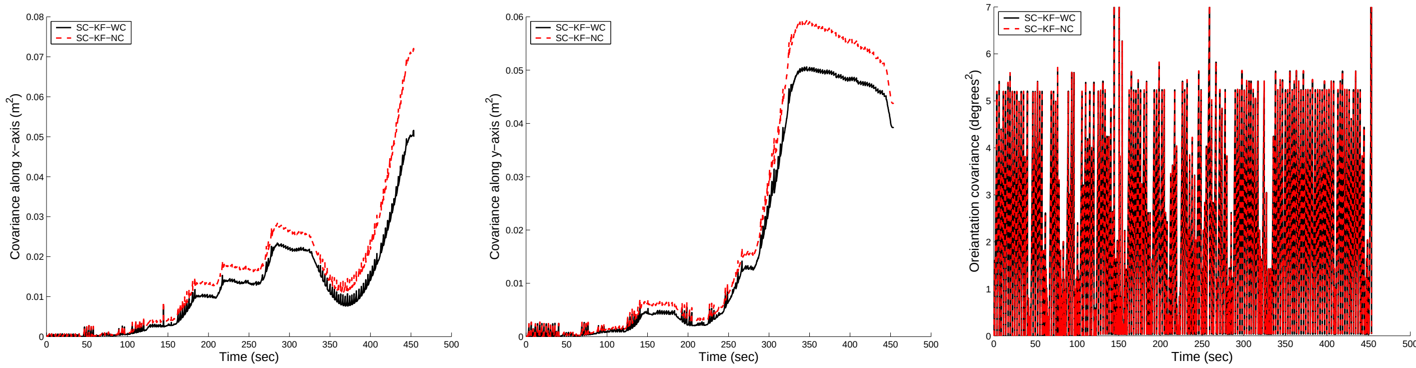

(a) (b) (c)

Fig. 2. The time evolution of the diagonal elements of the covariance matrix of the robot's pose. Note the difference in the vertical axes' scale. The intense fluctuations in the robot's orientation covariance arise due to the very high accuracy of the relative orientation measurements, compared to the low accuracy of the odometry-based orientation estimates. (a) covariance along the x-axis (b) covariance along the y-axis (c) orientation covariance.

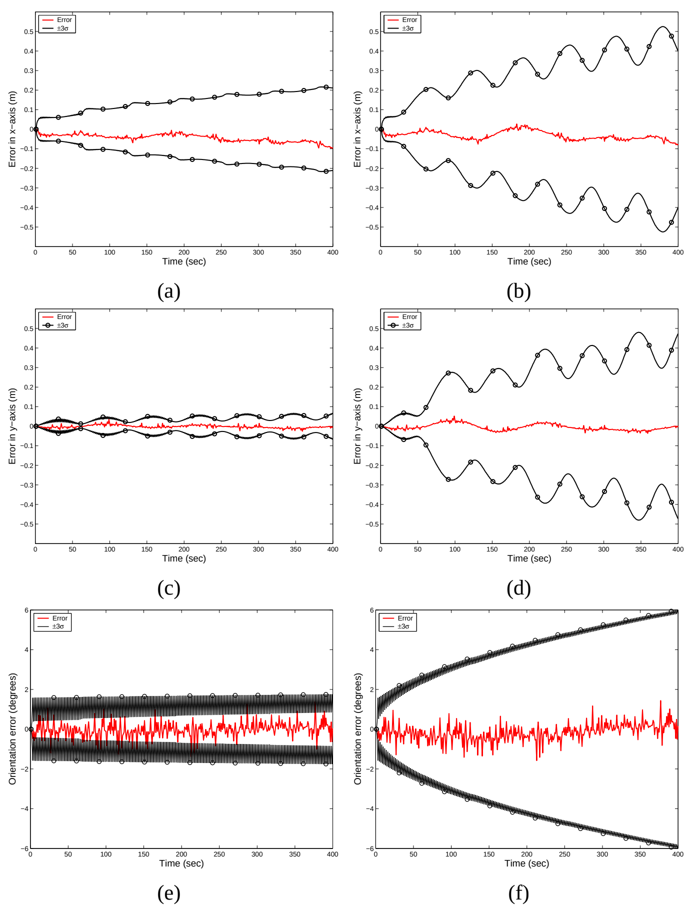

Fig. 3. The robot pose errors (solid lines) vs. the corresponding $99,8\%$ percentile of their distribution, (dashed lines with circles). The left column shows the results for the SC-KF-WC algorithm proposed in this paper, while the right one demonstrates the results for the SC-KF-NC algorithm. The "dark zones" in the last figures are the result of an intense sawtooth pattern in the robot's orientation variance. These fluctuations arise due to the very high accuracy of the relative orientation measurements, compared to the low accuracy of the odometry-based orientation estimates. (a - b) Errors along the x-axis (c - d) Errors along the y-axis (e - f) Orientation errors.

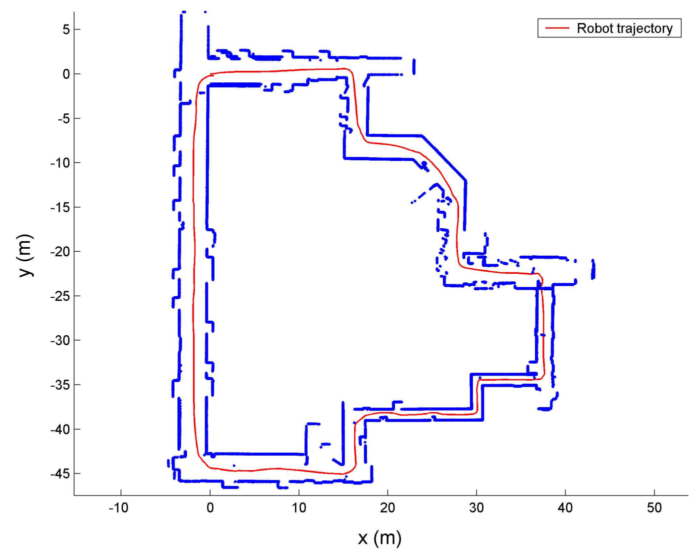

Fig. 4. The estimated trajectory of the robot based only on laser scan matching. The map is presented for visualization purposes only, by transforming all the laser points using the estimated robot pose.

in significantly *smaller* values for the covariance of the robot's pose estimates. Based on numerous experiments and simulation tests, it appears that this is a common result, which indicates that the correlation between consecutive displacement estimates tends to be *negative*.

An intuitive explanation for this observation can be given by means of a simple example, for motion on a straight line: as shown in Section III, the correlation of consecutive displacement estimates is attributed to the fact that the measurement errors at time step k affect the displacement estimates for both time intervals $[k-1,k]$ and $[k,k+1]$. Consider a robot moving along a straight-line path towards a feature, while measuring its distance to it at every time step. If at time-step k the error in the distance measurement is equal to $\epsilon_k > 0$, this error will contribute towards *underestimating* the robot's displacement during the interval $[k-1,k]$, but will contribute towards *overestimating* the displacement during the interval $[k,k+1]$. Therefore the error ϵ_k has *opposite* effects on the two

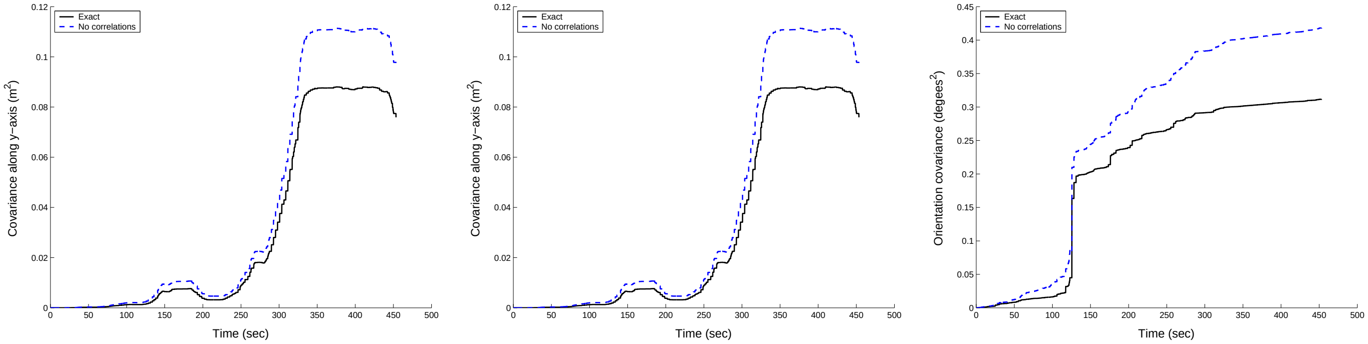

(a) (b) (c)

Fig. 5. The estimated covariance of the robot's pose using when the correlation between consecutive measurements is properly accounted for (solid lines) vs. the covariance estimated when the correlations are ignored (dashed lines). (a) Errors along the x-axis (b) Errors along the y-axis (c) Orientation errors.

displacement estimates, rendering them negatively correlated.

VI. Conclusions

In this paper, we have studied the problem of localization using relative-state measurements that are inferred from exteroceptive information. It has been shown, that when the same exteroceptive sensor measurements are employed for the computation of two consecutive displacement estimates (both forward and backward in time), these estimates are correlated. An analysis of the exact structure of the correlations has enabled us to derive an accurate formula for propagating the covariance of the pose estimates, which is applicable in scenarios when the exteroceptive measurements are the only available source of positioning information. To address the case in which proprioceptive sensor data are also available, we have proposed an efficient EKF-based estimation algorithm, that correctly accounts for the correlations attributed to the relative-state measurements. The experimental results demonstrate that the performance of the algorithm is superior to that of previous approaches [1], while the overhead imposed by the additional complexity is minimal. The method yields more accurate estimates, and most significantly, it provides a more precise description of the uncertainty in the robot's state estimates, thus facilitating motion planning and navigation.

Acknowledgements

This work was supported by the University of Minnesota (DTC), the Jet Propulsion Laboratory (Grant No. 1251073, 1260245, 1263201), and the National Science Foundation (ITR-0324864, MRI-0420836).

References

[1] S. I. Roumeliotis and J. Burdick, "Stochastic cloning: A generalized framework for processing relative state measurements," in *IEEE International Conference on Robotics and Automation*, Washington D.C., 2002, pp. 1788–1795.

[2] A. Kelly, "General solution for linearized systematic error propagation in vehicle odometry," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Maui, HI, Oct.29-Nov.3 2001, pp. 1938–45.

[3] L. Matthies, "Dynamic stereo vision," Ph.D. dissertation, Dept. of Computer Science, Carnegie Mellon University, 1989.

[4] H. S. C.F. Olson, L.H. Matthies and M. Maimone, "Robust stereo egomotion for long distance navigation," in *Proceedings of CVPR*, 2000, pp. 453–458.

[5] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2d range scans," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 18, no. 3, pp. 249–275, Mar. 1997.

[6] D. Silver, D. M. Bradley, and S. Thayer, "Scan matching for flooded subterranean voids," in *Proceedings of the 2004 IEEE Conf. on Robotics, Automation and Mechatronics (RAM)*, December 2004.

[7] S. T. Pfister, K. L. Kriechbaum, S. I. Roumeliotis, and J. W. Burdick, "Weighted range sensor matching algorithms for mobile robot displacement estimation," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Washington D.C., May 11-15 2002, pp. 1667–74.

[8] P. Torr and D. Murray, "The development and comparison of robust methods for estimating the fundamental matrix," *International Journal of Computer Vision*, vol. 24, no. 3, pp. 271–300, 1997.

[9] K. Konolige, "Large-scale map-making." in *AAAI National Conference on Artificial Intelligence*, San Jose, CA, jul 2004, pp. 457–463.

[10] B. D. Hoffman, E. T. Baumgartner, T. L. Huntsberger, and P. S. Shenker, "Improved state estimation in challenging terrain," *Autonomous Robots*, vol. 6, no. 2, pp. 113–130, April 1999.

[11] S. Julier and J. Uhlman, "A non-divergent estimation algorithm in the presence of unknown correlations," in *Proceedings of the American Control Conference*, vol. 4, 1997, pp. 2369 – 2373.

[12] R. C. Smith, M. Self, and P. Cheeseman, *Autonomous Robot Vehicles*. Springer-Verlag, 1990, ch. Estimating Uncertain Spatial Relationships in Robotics, pp. 167–193.

[13] M. Montemerlo, "Fastslam: A factored solution to the simultaneous localization and mapping problem with unknown data association," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, July 2003.

[14] S. Thrun, Y. Liu, D. Koller, A. Ng, Z. Ghahramani, and H. Durrant-Whyte, "Simultaneous localization and mapping with sparse extended information filters," *International Journal of Robotics Research*, vol. 23, no. 7-8, pp. 693–716, Aug. 2004.

[15] P. S. Maybeck, *Stochastic Models, Estimation, and Control*. Academic Press, 1979, vol. I.

[16] T. Bailey, "Constrained initialisation for bearing-only SLAM," in *Proceedigs of the IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, Sept. 2003, pp. 1966–1971.

[17] J. Leonard, R. Rikoski, P. Newman, and M. Bosse, "Mapping partially observable features from multiple uncertain vantage points," *The International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 943–975, 2002.

[18] J. Neira and J. D. Tardos, "Data association in stochastic mapping using the joint compatibility test," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 890–897, 2001.

[19] A. Howard, "Multi-robot mapping using manifold representations," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, New Orleans, LA, April 2004, pp. 4198–4203.

[20] S. Pfister, S. Roumeliotis, and J. Burdick, "Weighted line fitting algorithms for mobile robot map building and efficient data representation," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, Sep. 14-19 2003, pp. 1304–1311.