

Localization with Directional Coordinates

Charles Champagne Cossette, Mohammed Shalaby, David Saussié, James Richard Forbes

Abstract—A coordinate system is proposed that replaces the usual three-dimensional Cartesian x, y, z position coordinates, for use in robotic localization applications. Range, azimuth, and elevation measurement models become greatly simplified, and, unlike spherical coordinates, the proposed coordinates do not suffer from the same kinematic singularities and angle wrap-around. When compared to Cartesian coordinates, the proposed coordinate system results in a significantly enhanced ability to represent the true distribution of robot positions, ultimately leading to large improvements in state estimation consistency.

I. INTRODUCTION

This paper considers the problem of estimating the position of a robot, relative to some landmark or reference point, using range (distance), azimuth, and elevation (RAE) measurements. This problem has many applications in robotics, such as indoor localization or multi-robot relative position estimation. RAE measurements can be generated from a variety of sensors, such as radar used to track aircraft and spacecraft, or point cloud measurements obtained from LIDAR [1, Ch. 6.4.3]. Recently, ultra-wideband (UWB) radio has also been used to provide angle of arrival measurements [2], which is essentially an azimuth measurement, in addition to the usual distance measurements obtained between two UWB transceivers. However, single-antenna UWB angle of arrival measurements can have high variance [3].

A typical probabilistic estimation algorithm will use the familiar Cartesian x, y, z coordinates, with Gaussian estimators reporting the mean and covariance of the estimated distribution. However, as will be shown in this paper, Cartesian coordinates are generally a poor way of parameterizing the position of a robot when relying heavily on RAE measurements, especially when these measurements are very noisy, or if the robot position is highly uncertain. For example, the distribution of positions given a range measurement generally appears spherical, as shown in Figure 1, which greatly differs from a standard Gaussian ellipsoid. For this reason, this paper proposes an alternate coordinate system, referred to herein as *directional coordinates*. These coordinates are shown to provide a much more accurate representation of the position distribution, ultimately leading to improved state estimate consistency.

The idea of changing the coordinate system to better represent a distribution is not new. The “banana distribution”

*This work was supported by the FRQNT under grant 2018-PR-253646, the Canadian Foundation for Innovation JELF, the William Dawson Scholar Program, and the NSERC Discovery Grant Program.

C. C. Cossette, M. Shalaby, J.R. Forbes are with the Department of Mech. Engineering, McGill University. charles.cossette@mail.mcgill.ca, mohammed.shalaby@mail.mcgill.ca, james.richard.forbes@mcgill.ca

D. Saussié is with the Department of Electrical Engineering, Polytechnique Montréal. d.saussie@polymtl.ca

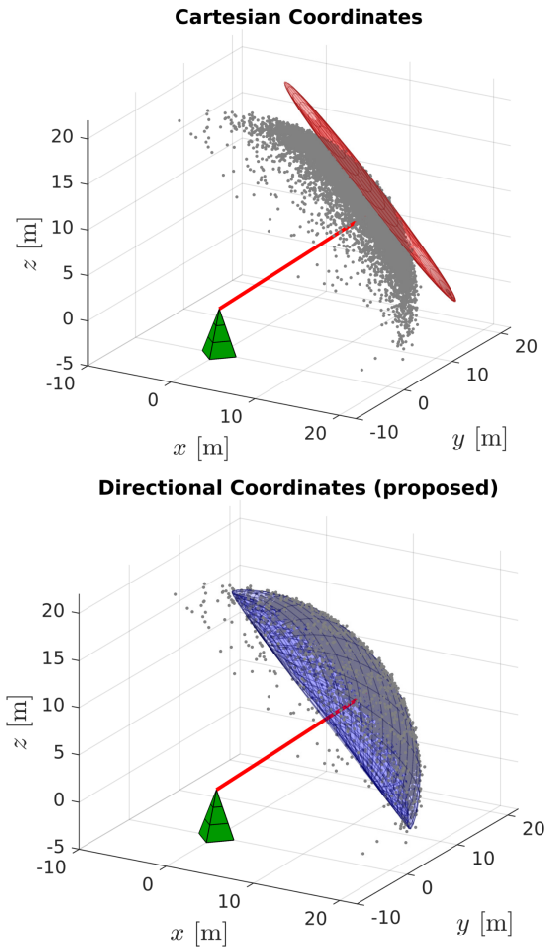


Fig. 1. Distribution of positions of a robot given a single distance measurement (red line) to a landmark (green pyramid) and a Gaussian prior distribution. The point cloud represents the “true” distribution, as determined by a particle filter, while the red and blue volumes are 3σ equal probability contours in Cartesian and directional coordinates, respectively.

associated with robot position uncertainty, primarily resulting from attitude uncertainty and sensor noise, has motivated the use of *exponential coordinates* [4]. A Gaussian distribution in exponential coordinates transforms to a “banana” distribution when mapped back to Cartesian coordinates. These exponential coordinates are naturally exploited when using matrix Lie groups in robotic estimation problems [1, 5]. Spherical coordinates use range, azimuth, and elevation as the coordinates themselves, and have been used in state estimation [6, 7]. Spherical coordinates are often modified to avoid singularities when the range is zero, such as using the logarithm of the range in [7], or its reciprocal [6]. However, spherical coordinates

possess an additional kinematic singularity at an elevation of $\varepsilon = \pm\pi/2$ rad, and also suffer from angle wrap-around, which must be overcome with explicit checks in the estimator implementation. Another strategy specific to RAE measurements is simply to directly convert them to a position measurement [8, 9], thus creating a linear measurement model. However, this requires careful conversion of the original noise statistics to equivalent noise statistics on the new measurement, which can introduce inaccuracies.

The new *directional coordinate* system proposed in this paper uses a range and direction cosine matrix (DCM) to parameterize position. The result is a smooth parameterization with 1) no angle wrap-around issues, 2) no kinematic singularities associated with the elevation, 3) RAE measurement models that are linear, or have constant, state-estimate-independent Jacobians, and 4) a filter that is significantly more consistent than a standard Cartesian coordinate filter, when noise levels are high. One drawback of the proposed use of directional coordinates is that a previously linear process model in Cartesian coordinates becomes nonlinear, an example being a simple “velocity input” process model. However, in applications with measurements arriving at a regular frequency, this added complexity is justified by the simplified measurement models, and enhanced ability to capture the true distributions. The proposed directional coordinates are evaluated in an extended Kalman filter (EKF) framework, and compared to an EKF that uses the usual Cartesian coordinates.

The paper is as follows. Directional coordinates are defined in Section II, with a basic probabilistic treatment shown in Section III. Section IV considers directional coordinate kinematics. Simulation and experimental results are shown in Sections V and VI, respectively.

A. Notation

A bolded $\mathbf{1}$ or $\mathbf{0}$ indicates an appropriately-sized identity or zero matrix, respectively. A Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ is written as $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The standard Euclidean 2-norm is denoted $\|\cdot\|$. The *special orthogonal group* of dimension 3 is denoted $SO(3)$ and whose elements belong to the set $\{\mathbf{C} \in \mathbb{R}^{3 \times 3} \mid \mathbf{C}^T \mathbf{C} = \mathbf{1}, \det \mathbf{C} = 1\}$. The operator $(\cdot)^\times : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ is the skew-symmetric cross-product matrix operator, with $\mathfrak{so}(3)$ being the Lie algebra of $SO(3)$. To reduce clutter, the argument of time (t) will be omitted from time-varying objects.

II. DIRECTIONAL COORDINATES

In this paper, *directional coordinates* are defined as the pair (ρ, \mathbf{C}) , where $\rho \in \mathbb{R}_{\geq 0}$ and $\mathbf{C} \in SO(3)$ is a direction cosine matrix (rotation matrix). The idea is to use these coordinates as an alternate representation of the components of a Cartesian position vector, denoted $\mathbf{r} \in \mathbb{R}^3$. This is done with the defining relation

$$\mathbf{r} = \rho \mathbf{C} \mathbf{e}_1, \quad (1)$$

where $\mathbf{e}_1 = [1 \ 0 \ 0]^T$. The *range* $\rho = \|\mathbf{r}\|$ is the length of the position vector, whereas \mathbf{C} can be interpreted as a direction

cosine matrix (DCM) that “rotates” the vector \mathbf{e}_1 to produce a *unit direction vector* $\mathbf{g} = \mathbf{C} \mathbf{e}_1 = \mathbf{r} / \|\mathbf{r}\|$ collinear to \mathbf{r} .

Directional coordinates are similar to spherical coordinates, but the azimuth and elevation angles are instead represented with a DCM \mathbf{C} . This avoids the kinematic singularities associated with spherical coordinates, as well as any angle wrap-around issues.

A. Cartesian to directional coordinates

Given a position vector $\mathbf{r} = [r_x \ r_y \ r_z]^T$ expressed in Cartesian coordinates, equivalent directional coordinates (ρ, \mathbf{C}) satisfying (1) can be obtained from the well-known axis-angle parameterization of a DCM. The range is straightforwardly obtained with $\rho = \|\mathbf{r}\|_2$, while the direction cosine matrix \mathbf{C} can be obtained by first defining \mathbf{a} and ψ as

$$\mathbf{a} = \frac{\rho \mathbf{e}_1^\times \mathbf{r}}{\|\rho \mathbf{e}_1^\times \mathbf{r}\|} = \frac{\begin{bmatrix} 0 \\ -r_z \\ r_y \end{bmatrix}}{\left\| \begin{bmatrix} 0 \\ -r_z \\ r_y \end{bmatrix} \right\|},$$

$$\psi = \arccos \frac{\rho \mathbf{e}_1^T \mathbf{r}}{\|\rho \mathbf{e}_1\| \|\mathbf{r}\|} = \arccos \frac{\rho \mathbf{e}_1^T \mathbf{r}}{\rho^2} = \arccos \frac{r_x}{\rho},$$

from which \mathbf{C} is then given by $\mathbf{C} = \exp(\psi \mathbf{a}^\times)$, which is well defined except when $r_z = r_y = 0$. In this case, setting $\rho = \|\mathbf{r}\|$ and $\mathbf{C} = \mathbf{1}$ will satisfy the definition given by (1). As such, a continuous function $\mathbf{h} : \mathbb{R}^3 \rightarrow \mathbb{R} \times SO(3)$ is defined that performs the above operations to go from \mathbf{r} to (ρ, \mathbf{C}) . That is,

$$(\rho, \mathbf{C}) = \mathbf{h}(\mathbf{r}). \quad (2)$$

B. Local parameterization

Estimation tools that operate on $SO(3)$ directly are very well established [1, 4, 5], most of which require the use of a *local parameterization* of $SO(3)$. In this paper, this is given by $\boldsymbol{\phi} = [\phi_1 \ \phi_2]^T \in \mathbb{R}^2$, defined such that

$$\mathbf{C} = \exp(\boldsymbol{\phi}^\wedge),$$

where the *wedge operator* $(\cdot)^\wedge : \mathbb{R}^2 \rightarrow \mathfrak{so}(3)$ is given by

$$\begin{bmatrix} \phi_1 \\ \phi_2 \end{bmatrix}^\wedge = \begin{bmatrix} 0 & -\phi_2 & \phi_1 \\ \phi_2 & 0 & 0 \\ -\phi_1 & 0 & 0 \end{bmatrix},$$

and otherwise expressed as $\boldsymbol{\phi}^\wedge = ([0 \ \boldsymbol{\phi}^T]^T)^\times$. An inverse mapping, named the *vee operator*, $(\cdot)^\vee : \mathfrak{so}(3) \rightarrow \mathbb{R}^2$ can be defined such that $((\boldsymbol{\phi}^\wedge)^\vee)^\wedge = \boldsymbol{\phi}$. It is also useful to define an operator $(\cdot)^\circ : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 2}$ such that the identity $\boldsymbol{\phi}^\wedge \mathbf{a} = \mathbf{a}^\circ \boldsymbol{\phi}$ holds. With the above definition of the wedge operator, it can be shown that the $(\cdot)^\circ$ operator must be

$$\mathbf{a}^\circ = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}^\circ = \begin{bmatrix} a_3 & -a_2 \\ 0 & a_1 \\ -a_1 & 0 \end{bmatrix}. \quad (3)$$

A few useful identities can be derived, for $\mathbf{b}, \mathbf{c} \in \mathbb{R}^2$, by simple expansion into components,

$$\mathbf{b}^{\wedge T} = -\mathbf{b}^\wedge, \quad (\mathbf{b} + \mathbf{c})^\wedge = \mathbf{b}^\wedge + \mathbf{c}^\wedge, \quad (\mathbf{b}^\wedge \mathbf{c}^\wedge)^\vee = \mathbf{0}. \quad (4)$$

III. DISTRIBUTIONS IN DIRECTIONAL COORDINATES

Consider the “mean” or “nominal” directional coordinates $(\check{\rho}, \check{\mathbf{C}})$ and the perturbations $\delta\rho$ and $\delta\phi$ such that

$$\rho = \check{\rho} + \delta\rho, \quad (5)$$

$$\mathbf{C} = \check{\mathbf{C}} \exp(\delta\phi^\wedge). \quad (6)$$

Define

$$\delta\mathbf{x} = \begin{bmatrix} \delta\rho \\ \delta\phi \end{bmatrix} \in \mathbb{R}^3. \quad (7)$$

If $\delta\mathbf{x}$ is a random variable drawn from $\delta\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \check{\mathbf{P}})$, this will induce a distribution over the values of (ρ, \mathbf{C}) .

A. Converting a Cartesian Gaussian distribution to directional coordinates

A common requirement will be to convert a Gaussian distribution over Cartesian coordinates $\mathbf{r} \sim \mathcal{N}(\check{\mathbf{r}}, \check{\mathbf{P}}_r)$, into an equivalent, or approximate, distribution in directional coordinates, represented with nominal point $(\check{\rho}, \check{\mathbf{C}})$ and corresponding distribution $\delta\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \check{\mathbf{P}})$. Fortunately, a mapping $\mathbf{h}(\cdot)$ has been identified in Section II-A, which allows for many well-known methods of passing a Gaussian through a nonlinearity [1, Ch. 2.2.8]. This paper uses sigma points for this step exclusively.

A variety of sigma point methods, such as the *unscented transform*, *spherical cubature*, or *Gauss-Hermite cubature* [10, Ch. 6], can be used to generate a set of sigma points and corresponding weights (\mathbf{s}_i, w_i) , $i = 1, \dots, N$ drawn from the prior distribution $\mathcal{N}(\check{\mathbf{r}}, \check{\mathbf{P}}_r)$. The Cartesian sigma points are transformed to directional coordinates through $\mathbf{h}(\cdot)$,

$$(\rho_i, \mathbf{C}_i) = \mathbf{h}(\mathbf{s}_i), \quad i = 1, \dots, N.$$

The covariance associated with the directional coordinates can then be approximated as

$$\check{\mathbf{P}} \approx \sum_{i=1}^N w_i \delta\xi_i \delta\xi_i^\top, \quad \delta\xi_i = \begin{bmatrix} \rho_i - \check{\rho} \\ \ln(\check{\mathbf{C}}^\top \mathbf{C}_i)^\vee \end{bmatrix}.$$

In standard sigma point methods, $(\check{\rho}, \check{\mathbf{C}})$ would be obtained from the weighted mean of the transformed sigma point




Fig. 2. Samples from a Cartesian Gaussian distribution, shown as the point cloud. The bean-shaped volume shows a 3 standard deviation equal probability contour of a Gaussian distribution in directional coordinates, but mapped back to Cartesian coordinates for visualization.

values (ρ_i, \mathbf{C}_i) [10, Ch. 6], which would require an optimization procedure to find the mean $\check{\mathbf{C}}$ on $SO(3)$ [11]. However, as suggested by [12], and verified extensively through simulation in this paper, it is sufficient to simply obtain the nominal points $(\check{\rho}, \check{\mathbf{C}})$ by passing $\check{\mathbf{r}}$ through the nonlinear model (2), $(\check{\rho}, \check{\mathbf{C}}) = \mathbf{h}(\check{\mathbf{r}})$, without an apparent loss in accuracy. An example of the results of this sigma point procedure can be visualized in Figure 2.

B. Posterior distribution given a range measurement

Consider now the task of estimating the distribution of (ρ, \mathbf{C}) given a prior distribution and a single range measurement y from the origin. The prior distribution’s nominal point is denoted $(\check{\rho}, \check{\mathbf{C}})$, and is distributed as per (5)-(6) with $\delta\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \check{\mathbf{P}})$. In directional coordinates, the range measurement model is linear,

$$y = \rho + \nu, \quad \nu \sim \mathcal{N}(0, R).$$

This leads to a trivial linearization procedure, where small changes in δy are related to small $\delta\mathbf{x}$ through $\delta y = \mathbf{H}\delta\mathbf{x} + \mathbf{M}\nu$, with $\mathbf{H} = [1 \ 0 \ 0]$ and $\mathbf{M} = 1$. This provides a setup to use a *multiplicative extended Kalman filter* [1] correction step to estimate the posterior mean $(\hat{\rho}, \hat{\mathbf{C}})$ and covariance $\hat{\mathbf{P}}$ with

$$\mathbf{K} = \check{\mathbf{P}}\mathbf{H}^\top(\mathbf{H}\check{\mathbf{P}}\mathbf{H}^\top + \mathbf{M}R\mathbf{M}^\top)^{-1}, \quad (8)$$

$$\delta\check{\mathbf{x}} = \begin{bmatrix} \delta\check{\rho} \\ \delta\check{\phi} \end{bmatrix} = \mathbf{K}\mathbf{z}, \quad \mathbf{z} = y - \check{\rho}, \quad (9)$$

$$\hat{\rho} = \check{\rho} + \delta\check{\rho}, \quad (10)$$

$$\hat{\mathbf{C}} = \check{\mathbf{C}} \exp(\delta\check{\phi}^\wedge), \quad (11)$$

$$\hat{\mathbf{P}} = (\mathbf{1} - \mathbf{K}\mathbf{H})\check{\mathbf{P}}(\mathbf{1} - \mathbf{K}\mathbf{H})^\top + \mathbf{K}\mathbf{M}R\mathbf{M}^\top\mathbf{K}^\top. \quad (12)$$

Figure 1 visualizes the results of this single EKF correction step when it is performed in Cartesian coordinates and directional coordinates, for the same Cartesian Gaussian prior. The point cloud is the distribution determined by a standard bootstrap particle filter [10, Ch. 7], and is considered the closest approximation to the true distribution. In Cartesian coordinates, the measurement model is given by $y = \|\mathbf{r}\|_2 + \nu$, which is nonlinear, and requires linearization to produce a state-dependent measurement

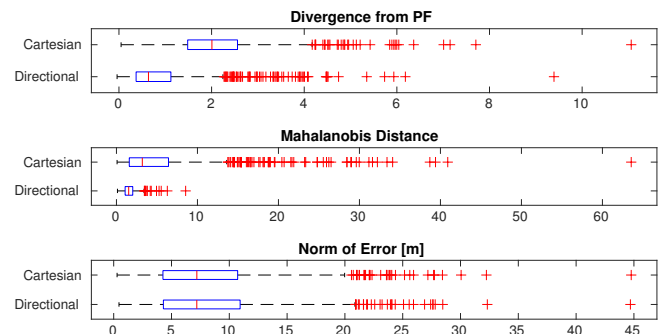


Fig. 3. 1000 Monte Carlo trials of a single Kalman correction given a distance measurement. Although the estimation accuracy using directional coordinates is not necessarily improved, directional coordinates offer a more consistent filter, as well as one that better resembles the particle filter.

Jacobian. Linearization errors aside, Figure 1 clearly shows the degree to which the distribution is non-Gaussian, when expressed in Cartesian coordinates. As such, under high uncertainty, there is no Cartesian Gaussian estimator that could ever accurately represent this distribution.

This simple single-step correction experiment is repeated for 1000 Monte Carlo trials, where the prior mean $\check{\mathbf{r}}$ and covariance $\check{\mathbf{P}}_r$ are randomized, and the true position \mathbf{r}^{true} is randomly sampled from this prior distribution. The prior distribution is converted to directional coordinates using Section III-A, and the posterior is calculated using Section III-B. Figure 3 displays various metrics of the 1000 Monte Carlo trials. The “dissimilarity”, or divergence, from the particle filter is calculated using the method in [13], which calculates an approximate Kullback-Leibler (KL) divergence between two point clouds, created by sampling their respective distributions. The Mahalanobis distance for the directional coordinates is calculated with

$$D = \sqrt{\delta \xi^T \hat{\mathbf{P}}^{-1} \delta \xi}, \quad \delta \xi = \begin{bmatrix} \rho^{\text{true}} - \hat{\rho} \\ \ln(\check{\mathbf{C}}^T \mathbf{C}^{\text{true}})^\vee \end{bmatrix}.$$

Although, on a single step, there is no accuracy improvement when using directional coordinates, there are substantially enhanced consistency properties, as reflected by the significantly lower Mahalanobis distances and divergence from the particle filter.

C. Posterior distribution given azimuth and elevation measurements

Another set of measurements that are naturally well-suited to directional coordinates are azimuth and elevation (AE) measurements, denoted α and ε respectively. Azimuth and elevation measurements are defined to produce a direction vector $\boldsymbol{\rho}$ where

$$\boldsymbol{\rho} = \mathbf{C}_z(\alpha)^T \mathbf{C}_y(-\varepsilon)^T \mathbf{e}_1, \quad (13)$$

and $\mathbf{C}_z(\cdot)$, $\mathbf{C}_y(\cdot)$ are principle rotation DCMs about the z and y axes, respectively. Hence, if the direction vector resulting from (13) is considered to be the measurement itself, the measurement model

$$\mathbf{y} = \mathbf{C}_{\mathbf{e}_1} + \boldsymbol{\nu}, \quad \boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}).$$

is obtained. The covariance \mathbf{R} can be obtained from covariances associated with α and ε using a standard linearization procedure [1, Ch. 2.2.8]. Next, a measurement innovation \mathbf{z} is defined following inspiration from the *invariant extended Kalman filter* (IEKF) literature [5]. Moreover, only two components of \mathbf{z} are used to avoid creating a fictitious third measurement, when only two are actually measured, that being α and ε . This is done through use of the projection matrix \mathbf{E} in (14), where the innovation

is defined as

$$\begin{aligned} \mathbf{z} &\triangleq \underbrace{\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{E}} \check{\mathbf{C}}^T (\mathbf{y} - \check{\mathbf{y}}) \\ &= \check{\mathbf{E}} \check{\mathbf{C}}^T \mathbf{C}_{\mathbf{e}_1} + \check{\mathbf{E}} \check{\mathbf{C}}^T \boldsymbol{\nu} - \check{\mathbf{E}} \check{\mathbf{C}}^T \check{\mathbf{C}}_{\mathbf{e}_1} \\ &\approx \mathbf{E}(\mathbf{1} + \delta \phi^\wedge) \mathbf{e}_1 + \check{\mathbf{E}} \check{\mathbf{C}}^T \boldsymbol{\nu} - \mathbf{E} \mathbf{e}_1 \\ &= \mathbf{E} \mathbf{e}_1^\odot \delta \phi + \check{\mathbf{E}} \check{\mathbf{C}}^T \boldsymbol{\nu} \triangleq \mathbf{H}_\phi \delta \phi + \mathbf{M} \boldsymbol{\nu}, \end{aligned} \quad (14)$$

where $\mathbf{C} = \check{\mathbf{C}} \exp(\delta \phi^\wedge)$, and the first-order approximation $\exp(\delta \phi^\wedge) \approx (\mathbf{1} + \delta \phi^\wedge)$ has been made. As with the range measurements, the measurement Jacobian $\mathbf{H} = [\mathbf{0} \ \mathbf{H}_\phi]$ is constant and state-estimate independent, and would otherwise be highly nonlinear in Cartesian coordinates. One may return to (8)-(12) to compute a Kalman filter correction step, with new definitions for \mathbf{y} , \mathbf{z} , \mathbf{H} , \mathbf{M} , and \mathbf{R} .

IV. DIRECTIONAL COORDINATE KINEMATICS

To execute dynamic filtering tasks, a process model is required, and it will usually be necessary to relate the rate of change of the directional coordinates ($\dot{\rho}$, $\dot{\mathbf{C}}$) to a Cartesian velocity input $\mathbf{v} = \dot{\mathbf{r}}$. Differentiating (1) with respect to time,

$$\dot{\mathbf{r}} = \dot{\rho} \mathbf{C}_{\mathbf{e}_1} + \rho \dot{\mathbf{C}}_{\mathbf{e}_1}. \quad (16)$$

The time rate of change of \mathbf{C} is $\dot{\mathbf{C}} = \mathbf{C} \boldsymbol{\omega}^\wedge$, where $\boldsymbol{\omega} \in \mathbb{R}^2$. Substituting $\dot{\mathbf{C}} = \mathbf{C} \boldsymbol{\omega}^\wedge$ into (16), and using the $(\cdot)^\odot$ operator defined in Section II-B results in

$$\begin{aligned} \dot{\mathbf{r}} &= \dot{\rho} \mathbf{C}_{\mathbf{e}_1} + \rho \mathbf{C} \boldsymbol{\omega}^\wedge \mathbf{e}_1 \\ &= \dot{\rho} \mathbf{C}_{\mathbf{e}_1} + \rho \mathbf{C} \mathbf{e}_1^\odot \boldsymbol{\omega} \\ &= \begin{bmatrix} \mathbf{C}_{\mathbf{e}_1} & \rho \mathbf{C} \mathbf{e}_1^\odot \end{bmatrix} \begin{bmatrix} \dot{\rho} \\ \boldsymbol{\omega} \end{bmatrix} \\ &= \underbrace{\mathbf{C} \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_1^\odot \end{bmatrix}}_{\mathbf{S}(\rho, \mathbf{C})} \begin{bmatrix} 1 & 0 \\ 0 & \rho \mathbf{1} \end{bmatrix} \begin{bmatrix} \dot{\rho} \\ \boldsymbol{\omega} \end{bmatrix}. \end{aligned} \quad (17)$$

To obtain the time rate of change of the directional coordinates, \mathbf{S} must be invertible. Fortunately, assuming $\rho > 0$, \mathbf{S} has an analytical inverse given by

$$\mathbf{S}(\rho, \mathbf{C})^{-1} = \begin{bmatrix} 1 & 0 \\ 0 & (1/\rho) \mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{e}_1^T \\ \mathbf{e}_1^{\odot T} \end{bmatrix} \mathbf{C}^T.$$

It follows that

$$\begin{bmatrix} \dot{\rho} \\ \boldsymbol{\omega} \end{bmatrix} = \mathbf{S}(\rho, \mathbf{C})^{-1} \mathbf{v}, \quad (18)$$

which is well defined provided $\rho > 0$. Equation (18) can be easily split into its two constituent equations, leading to the time rate of change of the directional coordinates, that being

$$\dot{\rho} = \mathbf{e}_1^T \mathbf{C}^T \mathbf{v}, \quad \dot{\mathbf{C}} = \frac{1}{\rho} \mathbf{C} (\mathbf{e}_1^{\odot T} \mathbf{C}^T \mathbf{v})^\wedge. \quad (19)$$

The only kinematic singularity is located at $\rho = 0$, as opposed to spherical coordinates, which have an additional singularity at an elevation of $\varepsilon = \pm \pi/2$ rad. This singularity is not of concern provided the robot device does not move near the origin. This is often the case in practice, as if another robot or measurement device is located at the origin, the hardware itself forces a minimum separating distance.




Fig. 4. Simulation results from a single trial. Black lines represent the $\pm 3\sigma$ confidence bounds. The total estimation error on the right figure is calculated for both filters with $\mathbf{e} = [\mathbf{r}^{\text{true}\top} \mathbf{v}^{\text{true}\top}]^T - [\mathbf{r}^T \mathbf{v}^T]^T$.




Fig. 5. Average estimation error and NEES, throughout time, for 100 Monte Carlo trials. The dashed line represents the one-sided 99.7% probability boundary.

A. Linearization

It is necessary to linearize the directional coordinate kinematics in order to execute an EKF prediction step. This is done by applying small perturbations to ρ , \mathbf{C} , \mathbf{v} with $\rho = \hat{\rho} + \delta\rho$, $\mathbf{C} \approx \hat{\mathbf{C}}(\mathbf{1} + \delta\phi^\wedge)$, and $\mathbf{v} = \hat{\mathbf{v}} + \delta\mathbf{v}$. Equation (19) becomes

$$\hat{\rho} + \delta\rho \approx \mathbf{e}_1^\top (\hat{\mathbf{C}}(\mathbf{1} + \delta\phi^\wedge))^\top (\hat{\mathbf{v}} + \delta\mathbf{v}). \quad (20)$$

Expanding (20), neglecting higher-order terms, and applying identities (4), eventually gives

$$\delta\rho \approx -\mathbf{e}_1^\top (\hat{\mathbf{C}}^\top \hat{\mathbf{v}})^\odot \delta\phi + \mathbf{e}_1^\top \hat{\mathbf{C}}^\top \delta\mathbf{v}.$$

Linearizing the DCM kinematics gives

$$\frac{d(\mathbf{1} + \delta\phi^\wedge)}{dt} \approx \frac{d(\hat{\mathbf{C}}^\top \mathbf{C})}{dt} = \dot{\hat{\mathbf{C}}}^\top \mathbf{C} + \hat{\mathbf{C}}^\top \dot{\mathbf{C}}, \quad (21)$$

$$\delta\dot{\phi}^\wedge \approx \frac{-1}{\hat{\rho}} (\mathbf{e}_1^\odot \hat{\mathbf{C}}^\top \hat{\mathbf{v}})^\wedge \hat{\mathbf{C}}^\top \dot{\mathbf{C}} (\mathbf{1} + \delta\phi^\wedge) + \frac{1}{\hat{\rho}} \hat{\mathbf{C}}^\top \mathbf{C} (\mathbf{e}_1^\odot \hat{\mathbf{C}}^\top \mathbf{v})^\wedge. \quad (22)$$

Again, after expansion, neglecting higher-order terms, applying identities (4), and application of the $(\cdot)^\vee$ to both sides results in

$$\delta\dot{\phi} \approx \frac{1}{\hat{\rho}} \mathbf{e}_1^\odot \hat{\mathbf{C}}^\top \delta\mathbf{v} - \frac{1}{\hat{\rho}} \mathbf{e}_1^\odot \hat{\mathbf{C}}^\top (\hat{\mathbf{C}}^\top \hat{\mathbf{v}})^\odot \delta\phi - \frac{1}{\hat{\rho}^2} \mathbf{e}_1^\odot \hat{\mathbf{C}}^\top \hat{\mathbf{v}} \delta\rho. \quad (23)$$

V. SIMULATION

Two different EKFs are tested in simulation. The only difference between the two EKFs is the coordinate system used for the position parameterization. The standard EKF uses Cartesian coordinates, whereas the *directional coordinate Kalman filter* (DCKF) uses directional coordinates. In these simulations, it is assumed that the acceleration $\mathbf{a} = \dot{\mathbf{v}} = \ddot{\mathbf{r}}$ of some point is measured with noise, along with noisy RAE measurements. The EKF state consists of $[\mathbf{r}^\top \mathbf{v}^\top]^\top$, while the DCKF state is $(\rho, \mathbf{C}, \mathbf{v})$, where $\mathbf{v} \in \mathbb{R}^3$ is a Cartesian velocity.

For the DCKF, the continuous-time linearized process model can be written as $\delta\dot{\mathbf{x}} = \mathbf{A}\delta\mathbf{x} + \mathbf{L}\delta\mathbf{a}$ where $\delta\mathbf{x} = [\delta\rho \ \delta\phi^\top \ \delta\mathbf{v}^\top]^\top$,

$$\mathbf{A} = \begin{bmatrix} 0 & -\mathbf{e}_1^\top (\hat{\mathbf{C}}^\top \hat{\mathbf{v}})^\odot & \mathbf{e}_1^\top \hat{\mathbf{C}}^\top \\ -\frac{1}{\hat{\rho}^2} \mathbf{e}_1^\odot \hat{\mathbf{C}}^\top \hat{\mathbf{v}} & -\frac{1}{\hat{\rho}} \mathbf{e}_1^\odot \hat{\mathbf{C}}^\top (\hat{\mathbf{C}}^\top \hat{\mathbf{v}})^\odot & \frac{1}{\hat{\rho}} \mathbf{e}_1^\odot \hat{\mathbf{C}}^\top \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (24)$$

and $\mathbf{L} = [0 \ \mathbf{0} \ \mathbf{1}]^\top$. The prediction step of the DCKF integrates (19) and $\dot{\mathbf{v}} = \mathbf{a}$ forward in time with simple Euler integration, but any method can be used. An equivalent discrete-time linearized model $\delta\mathbf{x}_k = \mathbf{A}_{k-1}\delta\mathbf{x}_{k-1} + \mathbf{w}_{k-1}$, $\mathbf{w}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{k-1})$ can be obtained for the DCKF using a standard discretization technique such as a zero-order-hold. The matrices \mathbf{A}_{k-1} and \mathbf{Q}_{k-1} are used to propagate the filter covariance forwards with the standard expression, $\hat{\mathbf{P}}_k = \mathbf{A}_{k-1}\hat{\mathbf{P}}_{k-1}\mathbf{A}_{k-1}^\top + \mathbf{Q}_{k-1}$. The correction step for the DCKF is performed using equations (8)-(12), with \mathbf{y} , \mathbf{z} , \mathbf{H} , \mathbf{M} , and \mathbf{R} matrices defined in Sections III-B for a range measurement, or III-C for azimuth-elevation measurements. The \mathbf{H} matrices are augmented with zeros due to the additional state \mathbf{v} .

Figure 4 shows an example run with no covariance tuning, displaying the error behavior and consistency of the proposed DCKF. Figure 5 contains results for 100 Monte Carlo trials with mild covariance tuning. The *normalized estimation error squared* (NEES) [14, Ch. 5.4] η_k^i is a consistency metric, calculated for the directional coordinates at time step k , on trial i , with

$$\eta_k^i = \delta\xi_k^\top \hat{\mathbf{P}}_k^{-1} \delta\xi_k, \quad \delta\xi_k = \begin{bmatrix} \rho_k^{\text{true}} - \hat{\rho}_k \\ \ln(\hat{\mathbf{C}}_k^\top \mathbf{C}_k^{\text{true}})^\vee \\ \mathbf{v}_k^{\text{true}} - \hat{\mathbf{v}}_k \end{bmatrix}. \quad (25)$$




Fig. 6. Pozyx UWB Developer Tag mounted to a Raspberry Pi 4B.

TABLE I
SIMULATION NOISE PROPERTIES

Specification	Value	Units
Range std. dev.	0.1	m
AE std. dev.	0.8	rad
Accel. std. dev.	0.1	m/s ²
Init. pos. std. dev.	5	m
Init. vel. std. dev.	3	m/s
Meas. frequency	10	Hz

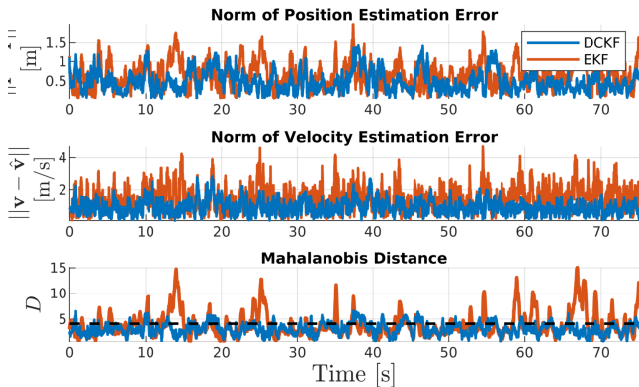


Fig. 7. Experimental results from a single trial, with the dashed line representing the one-sided 99.7% probability boundary.

Under large noise and initial uncertainty, the DCKF boasts an average 44% reduction in estimation error compared to the EKF and remains consistent, whereas the EKF is completely inconsistent, even with dedicated covariance tuning.

VI. EXPERIMENT

The EKF and DCKF are also compared in a real experiment. Figure 6 shows the Raspberry Pi 4B that was used with an LSM9DS1 IMU, providing accelerometer and gyroscope measurements, and a Pozyx UWB Developer Tag, providing distance measurements to a Pozyx UWB Creator Anchor on the floor. Because the particular UWB modules used do not provide azimuth and elevation measurements, azimuth and elevation measurements were simulated using ground-truth data, with artificially added zero-mean Gaussian noise with a standard deviation of 0.8 rad. This noise level is intentionally high, to induce a performance difference between the DCKF and EKF, and highlight the DCKF's improvement under large noise. The device moved in an overall volume of roughly $5 \text{ m} \times 4 \text{ m} \times 2 \text{ m}$. Ground truth position and attitude measurements are collected using an OptiTrack optical motion capture system.

Figure 7 shows various performance and consistency metrics associated with the experimental trial. The position and velocity root-mean-squared error (RMSE) was 0.55 m, 0.87 m/s for the DCKF, and 0.97 m, 1.56 m/s for the EKF, respectively. This corresponds to a 43% reduction in position RMSE, and a 44% reduction in velocity RMSE. Once again, the DCKF is both more accurate and more consistent, when compared to the EKF.

VII. CONCLUSION

This paper introduces a new coordinate system for parameterizing positions, and shows how it is particularly well-suited for RAE measurements. With these measurements, the new directional coordinate system captures the posterior distribution much more effectively than Cartesian coordinates, which culminates in a more accurate and significantly more consistent estimator. Future work could investigate eliminating the kinematic singularity at $\rho = 0$, by using the logarithm of the range as a coordinate, as done in [7].

REFERENCES

- [1] T. Barfoot, *State Estimation for Robotics*. Toronto, ON: Cambridge University Press, 2019.
- [2] I. Dotlic, A. Connell, H. Ma, J. Clancy, and M. McLaughlin, "Angle of Arrival Estimation Using Decawave DW1000 Integrated Circuits," *Workshop on Positioning, Navigation and Communications*, 2017.
- [3] A. Ledergerber, M. Hamer, and R. D'Andrea, "Angle of Arrival Estimation based on Channel Impulse Response Measurements," *IROS*, pp. 6686–6692, 2019.
- [4] A. W. Long, C. K. Wolfe, M. J. Mashner, and G. S. Chirikjian, "The Banana Distribution is Gaussian: A Localization Study with Exponential Coordinates," *Robotics: Science and Systems*, vol. 8, pp. 265–272, 2013.
- [5] A. Barrau and S. Bonnabel, "The Invariant Extended Kalman Filter as a Stable Observer," *IEEE Trans. on Automatic Control*, vol. 62, no. 4, pp. 1797–1812, 2017.
- [6] A. Boberg, A. N. Bishop, and P. Jensfelt, "Robocentric Mapping and Localization in Modified Spherical Coordinates with Bearing Measurements," *ISSNIP*, pp. 139–144, 2009.
- [7] M. Mallick, L. Mihaylova, S. Arulampalam, and Y. Yan, "Angle-only Filtering in 3D using Modified Spherical and Log Spherical Coordinates," in *Intl. Conf. on Information Fusion*, 2011.
- [8] D. Lerro and Y. Bar-Shalom, "Tracking With Debaised Consistent Converted Measurements Versus EKF," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 29, no. 3, pp. 1015–1022, 1993.
- [9] S. V. Bordonaro, "Converted Measurement Trackers for Systems with Nonlinear Measurement Functions," Ph.D. Thesis, University of Connecticut, 2015.
- [10] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge University Press, 2010, pp. 1–232.
- [11] Y. Dai, J. Trumpf, H. Li, N. Barnes, and R. Hartley, "Rotation Averaging with Application to Camera-Rig Calibration," in *ACCV*, 2009, pp. 335–346.
- [12] M. Brossard, A. Barrau, and S. Bonnabel, "A Code for Unscented Kalman Filtering on Manifolds (UKF-M)," in *ICRA*, 2020.
- [13] Q. Wang, S. R. Kulkarni, and S. Verdú, "Divergence Estimation for Multidimensional Densities via κ -Nearest-Neighbor Distances," *IEEE Trans. on Information Theory*, vol. 55, no. 5, pp. 2392–2405, 2009.
- [14] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley & Sons, Inc., 2001.