

Minimizing Bearing Bias in Tracking By De-coupled Rotation and Translation Estimates

Raman Arora
Department of Electrical Engineering
University of Washington
Seattle, WA 98195
Email: rmmnarora@u.washington.edu

Maya R. Gupta
Department of Electrical Engineering
University of Washington
Seattle, WA 98195
Email: gupta@ee.washington.edu

Abstract—The problem of tracking Euclidean motion is formulated as a sequential learning of rotations and translations. For tracking modalities such as radar and sonar, this approach avoids a fundamental mismatch that arises with standard trackers that model motion dynamics in Cartesian coordinates but track based on measurements whose noise is best modeled in polar coordinates. By considering motion in terms of rotations and translations and using group-theoretic estimation, the proposed tracker enjoys the advantage of unbiased averaging on the rotation group, in accordance with the geometry of the measurements. We demonstrate the proposed method with illustrative preliminary experiments. The stability and convergence of the proposed algorithm is established, extending known convergence results for online learning of rotations.

I. INTRODUCTION

The problem of tracking in applications like radar and sonar is a non-linear estimation problem since the target dynamics are modeled in Cartesian coordinates whereas the measurements are in spherical or polar coordinates. The classical Kalman filter is optimal only if the dynamical system as well as the observation model are linear, the state evolution is Markovian and the noise in observations is additive with Gaussian statistics [1].

For many practical sensing systems that take bearing and range measurements, the linearity assumption does not hold because of the nonlinear mapping between the measurements (in spherical coordinate system) and the states of the system (in Cartesian coordinates). Furthermore, the additive Gaussian modeling assumption is not valid because the noise in bearing (elevation and azimuth angles) measurements is on a 3D sphere which when transformed to Cartesian coordinates does not exhibit Gaussian statistics. To alleviate these problems, various extensions of the Kalman filter have been proposed that indirectly address these issues [2]–[4]. However, these methods still result in suboptimal performance due to inaccurate modeling and estimation of statistics of measurement noise. In particular, bearing measurement errors should be modeled as samples from a distribution on the unit sphere. Therefore, averaging them with respect to the Lebesgue measure on the Euclidean space (rather than the Haar measure on the rotation group) results in a biased estimator. Further, the noise in the two bearing angle measurements is typically assumed to be un-correlated, but this has not been confirmed

experimentally.

Any Euclidean motion can be specified in terms of a rotation and a translation. However, this representation is not unique. In this paper, we choose the rotation to represent differential bearing with respect to the receiver, and translation to represent the change in the range. This representation matches the geometry of noise for applications like sonar and radar and allows native modeling of the measurement noise statistics. Specifically, we model noise in bearing measurements (azimuth and elevation angles) using a spherical probability distribution (Kent distribution [5]), which can also incorporate correlation between bearing angles' noise. Based on this model, we present a sequential filtering approach which alternates between estimating rotation and translation.

The presented approach leverages recent advances in online learning of rotations [6]–[8]. Note that any motion in Euclidean space can be represented by a rotation and translation, in fact one can move from one point to another by multiple choices of a rotation and translation. We formulate the rotation to be differential bearing with respect to the receiver, and translation to be the change in the range, which matches the usual geometry of noise for applications like sonar and radar.

The paper is organized as follows. First, we discuss some related work in Section II. Then we overview our approach in Section III, and present the algorithm for simplified cases before presenting a general algorithm for tracking in Section IV. Finally, we present analysis of our algorithm in Section V, discuss experimental results in Section VI and conclude with an outline of future directions.

II. PRIOR WORK

The most relevant prior work is a state-of-the-art radar tracking method called *converted measurements Kalman filtering* (CMKF), which converts the observations from spherical coordinates to Cartesian coordinates and employs sequential non-linear filtering [4]. The main idea is to track the mean and covariance of the error of converted measurements in an extended Kalman filtering framework. Much of the recent research in this area has focused on making converted measurement error unbiased and improving estimates of covariance. Consider the following dynamic model for state-evolution:

$$X_{k+1} = \Phi_k X_k + G_k U_k + \Gamma_k W_k, \quad (1)$$

where X_k is the state vector with three components each for position, velocity, and acceleration, Φ_k is the state-transition matrix, G_k and Γ_k are coefficient or gain matrices, U_k is the control and W_k is additive observation noise with distribution $\mathcal{N}(0, Q_k)$. The observation model can be written as

$$Z_k^m = f(X_k) + V_k^m, \quad (2)$$

where $Z_k^m = [r_k^m \ \beta_k^m \ \theta_k^m \ \dot{r}_k^m]^T$ are the measurements, $f(X_k) = [r_k \ \beta_k \ \theta_k \ \dot{r}_k]^T$ are the true values and $V_k^m = [\tilde{r}_k \ \tilde{\beta}_k \ \tilde{\theta}_k \ \tilde{\dot{r}}_k]^T$ represents the noise. The function f is the Cartesian-to-spherical coordinate transformation map

$$\begin{aligned} r_k &= \sqrt{x_k^2 + y_k^2 + z_k^2}, \\ \dot{r}_k &= \frac{x_k \dot{x}_k + y_k \dot{y}_k + z_k \dot{z}_k}{r_k}, \\ \beta_k &= \tan^{-1} \left(\frac{z_k}{\sqrt{x_k^2 + y_k^2}} \right), \\ \theta_k &= \tan^{-1} \left(\frac{y_k}{x_k} \right). \end{aligned}$$

The converted position measurements x_k^c, y_k^c and z_k^c are given by spherical-to-Cartesian coordinate transform:

$$\begin{aligned} x_k^c &= r_k^m \cos(\beta_k^m) \cos(\theta_k^m), \\ y_k^c &= r_k^m \cos(\beta_k^m) \sin(\theta_k^m), \\ z_k^c &= r_k^m \sin(\beta_k^m). \end{aligned}$$

However, the conversion of Doppler measurements is highly non-linear and it is common to introduce a pseudo-measurement (denoted η_k^c) formed from the product of the Doppler and range measurements, $\eta_k^c = r_k^m \dot{r}_k^m = x_k \dot{x}_k + y_k \dot{y}_k + z_k \dot{z}_k + \tilde{\eta}_k$. In compact notation, the converted measurement can be written as,

$$Z_k^c = h_k(X_k) + V_k^c \quad (3)$$

where $Z_k^c = [x_k^c \ y_k^c \ z_k^c \ \eta_k^c]^T$ denotes converted measurements, $h_k(X_k) = [x_k \ y_k \ z_k \ (x_k \dot{x}_k + y_k \dot{y}_k + z_k \dot{z}_k)]^T$ are the true values and $V_k^c = [\tilde{x}_k \ \tilde{y}_k \ \tilde{z}_k \ \tilde{\eta}_k]^T$ is the additive noise in converted measurements. Kalman filtering then proceeds with estimated bias and variance of error of converted measurements. Denote the true bias and covariance of the converted measurement errors conditioned on the true position and range-rate of the target by $\mu_{k,t} = \mathbb{E}[V_k^c \mid r_k, \beta_k, \theta_k, \dot{r}_k]$, $R_{k,t} = \text{cov}(V_k^c, V_k^c \mid r_k, \beta_k, \theta_k, \dot{r}_k)$. Since the true values are not observed, in practice, the expected value of the true bias and covariance are evaluated conditioned on the measured position and range-rate: $\mu_{k,a} = \mathbb{E}[\mu_{k,t} \mid r_k^m, \beta_k^m, \theta_k^m, \dot{r}_k^m]$, $R_{k,a} = \mathbb{E}[R_{k,t} \mid r_k^m, \beta_k^m, \theta_k^m, \dot{r}_k^m]$.

Given the estimates $\mu_{k,a}$ and $R_{k,a}$, the position and pseudo-measurements are first de-correlated. Then estimation proceeds via sequential filtering where (a) target state is updated first using Kalman filter and converted position measurement and

(b) target state is updated via an extended Kalman filter (EKF) using pseudo-measurement.

There are several shortcomings of the approach described above. Firstly, the measurement errors in range, bearing and Doppler are typically assumed un-correlated although there is evidence of statistical correlations between bearing angles and between range and Doppler measurements. Secondly, the Gaussian statistics for measurement errors are not manifested as Gaussian in converted measurement errors. Finally, inconsistent statistics are obtained by the linearization of pseudo-measurements when using EKF.

III. PROPOSED APPROACH

We formulate the problem of tracking as sequential estimation of rotations and translations. The rotations are estimated using an approach designed for online-learning of rotations [7], [8]. We observe that a suitable choice of step-size sequence for the online algorithm results in averaging over the rotation group. This handles the noise in bearing measurements by averaging it on the sphere. Correlations between errors in bearing measurements are modeled directly on the unit sphere using the Kent distribution. The translations are tracked using a gradient descent algorithm.

Although our approach is quite general and extends to special Euclidean motion group $\text{SE}(n)$, we focus on applications like radar and SONAR for 3D tracking, and we take the receiver location to be the origin of the space.

Any Euclidean motion can be described as a composition of a rotation and translation, that is, any two points $p_1, p_2 \in \mathbb{R}^n$, are related via the linear transformation

$$p_2 = R p_1 + t, \quad (4)$$

where R is a $n \times n$ rotation matrix and $t \in \mathbb{R}^n$. The pair (R, t) for a given pair of points is not unique. We choose rotation to represent differential bearing with respect to the receiver, and translation to represent the change in the range. This matches the geometry of noise for applications like sonar and radar.

In the next subsections we begin with a simplified motion model of constant rotation and no translation to introduce the ideas and notation (Case I). We then extend to the case of changing rotation but no translation (Case II). Finally, we treat the practical case of changing rotations and translations (Case III). The notation we use in the rest of the paper is summarized in Table I.

A. Case I: Constant Rotation, No Translation

We begin with a simplified problem: learning a constant rotation from a series of bearing measurements. This scenario corresponds to a constant-curvature turn-maneuver of the target with respect to the observer, such as tracking the motion of an airplane flying at a constant altitude. Figure 1 shows an example of motion, measurements, and our estimates for this case.

Consider a sequence of state-vectors $x_n \in \mathbb{R}^3$ on the unit sphere S^2 that evolve as

$$x_n = R_* x_{n-1}, \quad (5)$$

Table I: Notation

Symbols	Description
$x_n \in \mathbb{R}^3$	state variable corresponding to bearing
$y_n \in \mathbb{R}^3$	bearing observations expressed in Cartesian coordinates
$\mathbf{SO}(3)$	3D rotation group
$R_n \in \mathbf{SO}(3)$	state corresponding to differential rotation of x_n
$t_n \in \mathbb{R}$	state corresponding to range
$s_n \in \mathbb{R}$	range observations
$\dot{t}_n \in \mathbb{R}$	state corresponding to the Doppler or range-rate
$\dot{s}_n \in \mathbb{R}$	Doppler or range-rate measurement
$(\hat{x}_n, \hat{R}_n, \hat{t}_n, \hat{\dot{t}}_n)$	estimates of the state vector $(x_n, R_n, t_n, \dot{t}_n)$
$z_n \in \mathbb{R}^3$	position of target in Cartesian coordinates
$\hat{z}_n \in \mathbb{R}^3$	estimated position of target in Cartesian coordinates
(α, β, γ)	Euler angles
(η_R, η_t, η_i)	step-sizes




Figure 1: Example Simulation of Case I: object moving on sphere of radius 2 units (no translation) from observer at origin. True motion is shown in green; the noisy bearing measurements are shown in blue; and the track estimated from these noisy measurements using proposed updates is shown in red.

where R_* is the unknown true rotation matrix that describes the constant rotation of the target about the receiver over a sampling interval. We assume that the sampling interval is a known constant. The observations are given as bearing measurements of the moving target and are denoted in Cartesian coordinates by $y_n \in \mathbb{R}^3$. We assume the following noise model for our observations:

$$y_n \sim \text{Kent}(x_n, \gamma_2, \gamma_3, \kappa, \beta), \quad (6)$$

i.e. y_n is sampled from a Kent distribution [5] with mean equal to the true bearing x_n of the target. The parameter κ controls the concentration of the distribution about the mean (larger the κ , more concentrated the distribution) and β controls the

ellipticity of equal probability contour ($\gamma_2, \gamma_3 \in \mathbb{R}^3$ determine the major and minor axes of the elliptical equal probability contours). This has the very practical advantage of allowing one to model the larger uncertainty in bearing angles in the direction of the motion, and the smaller uncertainty in the direction orthogonal to the motion. Note that we represent states and measurements corresponding to the bearing of the target in Cartesian coordinates.

This problem is similar to the problem of online learning of rotations [7], [8] where at each time instance one receives a vector on the unit sphere and its rotated version. In the tracking problem, the sequence of unit vectors (corresponding to the true bearing of the target) forms a Markov sequence, where the unit vector at instance n is the obtained by a rotation of the unit vector at instance $n-1$. We employ gradient-based multiplicative updates [7] for learning the rotation matrix.

Given the observation y_n at instance n , and an estimate \hat{R}_{n-1} of the constant rotation matrix, the state estimate \hat{x}_n is updated as:

$$\hat{x}_n = \hat{R}_{n-1} \hat{x}_{n-1}, \quad (7)$$

and the estimate of the rotation matrix is updated as

$$\hat{R}_n = \hat{R}_{n-1} \exp\left(-2\eta \text{skew}\left(\hat{R}_{n-1}^T (\hat{x}_n - y_n) \hat{x}_{n-1}^T\right)\right), \quad (8)$$

where $\exp(\cdot)$ denotes matrix exponential and $\text{skew}(\cdot)$ is defined as $\text{skew}(A) = A - A^T$ for any square matrix A . We show in Section V-A that these gradient updates with a suitable choice of step-size (see [7]) correctly average out zero-mean bearing noise. Convergence analysis for learning rotations is given in Theorem 2 in Section V-C.

B. Case II: Changing Rotation, No Translation

Next we extend our formulation to the scenario where the target is accelerating or slowing down while navigating a constant-curvature turn about the observer. The tracking problem in this case reduces to tracking rotations. We assume the following dynamical model:

$$\begin{aligned} x_n &= R_{n-1} x_{n-1}, \\ R_n &= f(R_{n-1}, \dot{R}_{n-1}), \end{aligned} \quad (9)$$

where we have added a new state R_n that captures the first order dynamics of x_n . The function $f(R_{n-1}, \dot{R}_{n-1})$ models the change in the rotation matrix (we assume the rotation matrix changes smoothly such that the evolution of the rotation matrix can be captured with the first order derivative of the rotation matrix). For instance, $f(R_{n-1}, \dot{R}_{n-1}) = R_{n-1} \exp(\eta \dot{R}_{n-1})$ describes an accelerating target that is turning in the direction given by \dot{R}_{n-1} at a rate controlled by η . Note that \dot{R}_{n-1} should be a skew-symmetric matrix as explained below. We will show that the update given in (8) is also optimal for this case.

We model the trajectory between two sequential bearing measurements y_{n-1} and y_n as a geodesic (shortest path on the sphere) which can be described as an element, say $A(\tau)$, of a continuous subgroup of rotation matrices parameterized

by $\tau \geq 0$, with extreme point $A(\tau) = I$ for $\tau = 0$. Recall that for any rotation matrix R , it holds that $RR^T = I$. Thus $A(\tau)A(\tau)^T = \mathbf{I}$ is a constant function of τ . Hence its derivative is zero, that is:

$$\frac{d(A(\tau)A(\tau)^T)}{d\tau} = A'(\tau)A(\tau)^T + A(\tau)A'(\tau)^T = 0. \quad (10)$$

In the limit as $\tau \rightarrow 0$, (10) reduces to $A'(\tau) + A'(\tau)^T = 0$, which implies that $A'(\tau)$ is skew-symmetric. Therefore, the change of the rotation matrix can be described by an element from the Lie algebra of the rotation group which describes the tangent direction for the change of the rotation matrix. This matches the evolution model assumed implicitly in the derivation of the updates in eqn. (8). Consequently, the online updates are optimal for tracking rotations given a suitable choice of step-size sequence (see Sec. V-A).

Let us consider these tracking estimates further, but specialize the discussion to 3D so that we can employ the Euler ZYZ parametrization of the rotation group. Any 3×3 rotation matrix R can be written as $R = R_Z(\alpha) R_Y(\beta) R_Z(\gamma)$, where $\alpha, \gamma \in [0, 2\pi]$, $\beta \in [0, \pi]$, $R_Z(\alpha)$ is the 3×3 rotation matrix corresponding to rotation about Z -axis by an angle α and $R_Y(\beta)$ is the 3×3 rotation matrix corresponding to rotation about Y -axis by an angle β . The dynamic model in terms of the Euler parameters can be written as $\alpha_n = \alpha_{n-1} + \eta_\alpha \dot{\alpha}_{n-1}$, $\beta_n = \beta_{n-1} + \eta_\beta \dot{\beta}_{n-1}$, $\gamma_n = \gamma_{n-1} + \eta_\gamma \dot{\gamma}_{n-1}$. Let $\theta_n = [\alpha_n, \beta_n, \gamma_n]^T$. In this compact notation, $\theta_n = \theta_{n-1} + \Psi \dot{\theta}_{n-1}$ where Ψ is a diagonal matrix with elements $[\eta_\alpha, \eta_\beta, \eta_\gamma]$. Also write $\dot{\theta}_n = \Phi \dot{\theta}_{n-1}$ where Φ models the angular acceleration of the target; $\Phi = \mathbf{I}_3$ models constant angular velocity. Then, matrix-exponentiated gradient updates in (8) track the following transition model,

$$\begin{bmatrix} \theta_n \\ \dot{\theta}_n \end{bmatrix} = \begin{bmatrix} \mathbf{I}_3 & \Psi \\ 0 & \Phi \end{bmatrix} \begin{bmatrix} \theta_{n-1} \\ \dot{\theta}_{n-1} \end{bmatrix}, \quad (11)$$

where the observations are sampled from a Kent distribution as described in (6).

C. Case III: Changing Rotations and Translation

Finally, we consider the practical scenario where movement is unrestricted. We add another state t_n that represents range of the target; the corresponding observation is denoted as s_n . The state transitions are:

$$\begin{aligned} x_n &= R_{n-1}x_{n-1}, \\ R_n &= f(R_{n-1}, \dot{R}_{n-1}), \\ t_n &= t_{n-1} + \eta_t \dot{t}_{n-1}, \\ \dot{t}_n &= \eta_i \dot{t}_{n-1}, \end{aligned} \quad (12)$$

where \dot{t}_n denotes the derivative of t_n . We model the bearing measurement y_n and range s_n and Doppler \dot{s}_n measurements as:

$$\begin{aligned} y_n &\sim \text{Kent}(x_n, \gamma_2, \gamma_3, \kappa, \beta), \\ \begin{pmatrix} s_n \\ \dot{s}_n \end{pmatrix} &\sim \mathcal{N}\left(\begin{pmatrix} t_n \\ \dot{t}_n \end{pmatrix}, \mathbf{Q}\right), \end{aligned} \quad (13)$$

where $\mathcal{N}(t, \mathbf{Q})$ is the bivariate Gaussian distribution with mean t and covariance \mathbf{Q} .

We propose an algorithm in Section IV to track the states. Convergence analysis of the proposed algorithm for learning Euclidean motion is given in Theorem 3 in Section V-D.

D. Decoupled State Vector

It is common in tracking for the state vector to be the Cartesian-coordinate position of the tracked object. Instead, we are proposing (see Sec. III-C) to independently track the bearing by learning rotations and range by learning translations. The state transitions for the bearing and range given by (12) are de-coupled. For example, for the three-dimensional case we parameterize the rotation matrix by the Euler angles θ and (12) can be written in a block-matrix form that makes the de-coupling explicit:

$$\begin{bmatrix} \theta_n \\ \dot{\theta}_n \\ t_n \\ \dot{t}_n \end{bmatrix} = \begin{bmatrix} \mathbf{I}_3 & \Psi & 0 & 0 \\ 0 & \Phi & 0 & 0 \\ 0 & 0 & 1 & \eta_t \\ 0 & 0 & 0 & \eta_i \end{bmatrix} \begin{bmatrix} \theta_{n-1} \\ \dot{\theta}_{n-1} \\ t_{n-1} \\ \dot{t}_{n-1} \end{bmatrix},$$

where Ψ and Φ are appropriate transition matrices for the Euler angles.

Consider an alternate formulation of the tracking problem in Case III where the state vector is chosen to be the position z_n of the tracked object in Cartesian coordinates. Then the state-transition in terms of a rotation and a translation would be:

$$z_n = \tilde{R}_{n-1}z_{n-1} + \tilde{t}_{n-1}, \quad (14)$$

where \tilde{R}_{n-1} represents the incremental rotation, about the receiver, relative to the last position z_{n-1} , and \tilde{t}_{n-1} represents the incremental radially outwards translation. At each step we would estimate the incremental rotation \hat{R}_{n-1} and \hat{t}_{n-1} and update the state estimate $\hat{z}_n = \hat{R}_{n-1}\hat{z}_{n-1} + \hat{t}_{n-1}$. We argue that this formulation, which uses the Cartesian coordinates z as the state vector, is not a good choice for tracking Euclidean motion with bearing and range measurements because it leads to biased estimates. To see this, treat the estimated rotation and translation as random (due to the randomness from the measurement noise) and consider the expectation of the estimate \hat{Z}_{n+1} after two time steps assuming one knew the earlier true position z_{n-1} :

$$\begin{aligned} & \mathbb{E} \left[\hat{Z}_{n+1} | \hat{Z}_{n-1} = z_{n-1} \right] \\ &= \mathbb{E} \left[\hat{R}_n \hat{R}_{n-1} z_{n-1} + \hat{R}_n \hat{t}_{n-1} + \hat{t}_n \right] \\ &= \underbrace{\mathbb{E} \left[\hat{R}_n \hat{R}_{n-1} \right]}_{\text{rotation}} z_{n-1} + \underbrace{\mathbb{E} \left[\hat{R}_n \hat{t}_{n-1} \right]}_{\text{coupling drift}} + \underbrace{\mathbb{E} \left[\hat{t}_n \right]}_{\text{translation}} \\ &= z_{n+1} + \left(\mathbb{E} \left[\hat{R}_n \hat{t}_{n-1} \right] - \tilde{R}_n \tilde{t}_{n-1} \right), \end{aligned} \quad (15)$$

where (16) follows from (15) if the estimates of the rotations and translation are unbiased. However, (16) shows that even with unbiased estimates of the rotation and translation,

the state vector \hat{Z}_{n+1} is biased by the coupling drift term. This bias problem motivates our proposed approach given in Section III-C, which differs from this alternate formulation in that we treat the bearing and range of the target as the states rather than the Euclidean position and then learn the differential rotation (with respect to the current bearing) and incremental translation (with respect to the current range).

IV. PROPOSED ALGORITHM

We propose tracking the Euclidean motion by alternating between estimating rotations and translations using a gradient descent method. Suppose one has the estimate $(\hat{R}_{n-1}, \hat{t}_{n-1})$ before contact n . The empirical errors can be expressed as $L_R = \|y_n - \hat{R}_{n-1}\hat{x}_{n-1}\|$, $L_t = \|s_n - \hat{t}_n\|_2^2$.

We form new estimates for the rotation and translation by taking a small step in the direction of the corresponding gradients $\nabla_R L_R$ and $\nabla_t L_t$. The gradient update for rotations is given by (8); we use the alternate (and equivalent) form as given in Theorem 1 in Section V. For translations, the gradient update is simply: $\hat{t}_n = \hat{t}_{n-1} + \eta_t \nabla_t L_t$. The resulting algorithm is given in Algorithm 1.

Input: Old estimates $(\hat{x}_{n-1}, \hat{R}_{n-1}, \hat{t}_{n-1}, \hat{t}_{n-1})$ of bearing, differential rotation, range and Doppler, step sizes (η_R, η_t, η_i) , sampling rate $\frac{1}{\Delta}$, new observations of bearing y_n , range s_n , and Doppler \dot{s}_n

Output: New estimates $(\hat{x}_n, \hat{R}_n, \hat{t}_n, \hat{t}_n)$ of bearing, differential rotation, range and Doppler

Predicted Bearing: $\hat{x}_n = \hat{R}_{n-1}\hat{x}_{n-1}$

Compute Matrix:

$$S = 2\eta \left(\hat{R}_{n-1}^T y_n x_{n-1}^T - x_{n-1} y_n^T \hat{R}_{n-1} \right)$$

Compute eigenvalue of S : $\lambda = 2\eta \sqrt{1 - (y_n^T \hat{x}_n)^2}$

New Rotation Estimate:

$$\hat{R}_n = \hat{R}_{n-1} \left(\mathbf{I} + \frac{\sin(\lambda)}{\lambda} S + \frac{1 - \cos(\lambda)}{\lambda^2} S^2 \right)$$

Predicted Translation: $\hat{t}_n = \hat{t}_{n-1} + \frac{1}{\Delta} \hat{t}_{n-1}$

Compute Gradient: $\nabla_t L_t = -2(s_n - \hat{t}_n)$

New Doppler Estimate:

$$\hat{t}_n = \hat{t}_{n-1} + \eta_t \nabla_t L_t + \eta_i (s_n - \hat{t}_n)$$

Algorithm 1: Tracking Updates at Instance n

V. ALGORITHM ANALYSIS

In this section we discuss convergence and stability results for Algorithm 1.

A. Proposed updates average bearing noise

The updates for learning rotations form a weighted averaging on the rotation group, with weights given by the step-sizes. In fact, they are a special case of a fixed-point iteration for computing mean on an arbitrary Lie group [9]. Consider a collection of i.i.d. samples $\{g_1, \dots, g_{n-1}\}$ from a probability distribution on the Lie group G with mean g . Many matrix Lie groups do not form a linear vector space; in other words not all matrix groups are closed under linear averaging, rotation group for instance. Consequently, linear combination $\frac{1}{n} \sum_k g_k$ is not necessarily an element in the group. Fortunately, each Lie group is associated with a Lie algebra which is the tangent space to the group at the identity element ($e \in G$) and admits the structure of a vector space. This inspires the following recipe for averaging on the Lie group: (a) transport each g_k to $\delta_k = g^{-1} \circ g_k$ in the neighborhood of e , (b) take the weighted average of the Lie algebra elements associated with δ_k , $\mu_g = \sum_k \eta_k \log(\delta_k)$, (c) retract back to the Lie group: $\bar{g} = \exp(\mu_g)$. Thus one needs weights η_k and element g that satisfy the relationship, $g = \bar{g}$, i.e.

$$g = g \circ \exp \left(\sum_{k=1}^{n-1} \eta_k \log(g^{-1} \circ g_k) \right). \quad (17)$$

For $G = R$, the group operation is commutative and the $\log(\cdot)$ and $\exp(\cdot)$ maps are identity maps. Thus, the equation above can be solved for η_k in a closed form giving $\eta_k = \frac{1}{n-1}$ for $k = 1, \dots, n-1$ and $\bar{g} = \frac{1}{n} \sum_k g_k$. For $G = \mathbf{SO}(3)$, however, the group operation is non-commutative and the retraction and the lifting map are matrix exponential and matrix logarithm respectively. Clearly, in this case, we may not solve for step sizes or the mean in a closed form but the following fixed-point iteration converges to the true mean for a suitable choice of step sizes,

$$g_n = g_{n-1} \circ \exp \left(\sum_{k=1}^{n-1} \eta_k \log(g_{n-1}^{-1} \circ g_k) \right). \quad (18)$$

For the rotation group the equation above reduces to updates in eqn. (8). Thus the learning-rotations updates yield an unbiased estimator on the rotation group.

B. Simplified updates

The updates presented in [7] can be written in two equivalent forms. While the matrix-exponentiated-gradient (see eqn. 8) form of the update is useful for geometrical intuition of the algorithm and its interpretation as moving averaging filter, the following simplified updates from [7] are useful for convergence analysis of the algorithm.

Theorem 1. (Complexity reduction [7]) *Let $S = 2\eta \left(\hat{R}_{n-1}^T y_n x_{n-1}^T - x_{n-1} y_n^T \hat{R}_{n-1} \right)$ be the skew-symmetric matrix in (8). The eigenvalues of S are given as $\pm j\lambda$ where $\lambda = 2\eta \sqrt{1 - (y_n^T \hat{x}_n)^2}$. Then the rotation estimates in (8) can be written equivalently as*

$$\hat{R}_n = \hat{R}_{n-1} \left(\mathbf{I} + \frac{\sin(\lambda)}{\lambda} S + \frac{1 - \cos(\lambda)}{\lambda^2} S^2 \right). \quad (19)$$

C. Convergence result for learning rotations

Assume the set-up of Case I, where there is one constant rotation R_* that we wish to estimate. The convergence of the updates given by (8) in the noiseless setting is given by the following theorem from [10]. Note that $\|\cdot\|_F$ denotes the Frobenius norm, and $\|\cdot\|_2$ denotes the l_2 -norm.

Theorem 2 (Stability of rotation estimates [10]). *The function $V(R) = \|R - R_*\|_F$, is non-increasing for updates in (8), i.e. $V(\hat{R}_n) \leq V(\hat{R}_{n-1})$ for all n .*

D. Convergence For Euclidean Motion Estimations

Next, consider the set-up of Case III, and assume the special case that the target is moving with a constant velocity. We establish here that the algorithm proposed in Algorithm 1 converges to the true Euclidean motion in the noiseless case. Theorem 3 follows as a corollary to Theorem 2.

Theorem 3 (Stability of Estimating Euclidean Motion). *The function $V(R, t) = \|R - R_*\|_F + \|t - t_*\|_2$ is non-increasing for Algorithm 1, i.e. $V(\hat{R}_{n+1}, \hat{t}_{n+1}) \leq V(\hat{R}_n, \hat{t}_n)$.*

VI. ILLUSTRATIVE EXPERIMENTAL RESULTS

This section presents some preliminary illustrative experimental results for the proposed tracking algorithm with simulated radar data. As a first experiment, consider a target moving at a constant velocity in a circular motion around the observer (Case I). A representative example is shown in Fig. 1 along with noisy measurements and the predicted track. The object is moving at a constant distance (2 units) from the observer located at origin. Fig. 2 shows the tracking performance of the proposed algorithm averaged over 1000 realizations for different levels of Kent noise specified by parameters κ and β . Note that the noise level in the bearing measurements decreases as κ increases. For each experiment in this section, we chose four different parameter values for Kent distribution: $(\kappa = 700, \beta = 300)$, $(\kappa = 500, \beta = 200)$, $(\kappa = 300, \beta = 100)$ and $(\kappa = 100, \beta = 0)$; the observation rate was chosen to be 20 samples per second and the target was tracked for 2 seconds. The tracking performance is evaluated in terms of the Euclidean distance between the actual and the predicted position of the object in \mathbb{R}^3 .

For a second experiment, we simulate a speeding target moving in a circular motion around the observer (Case II). Fig. 3 shows the tracking performance of the proposed algorithm averaged over 200 realizations for different levels of Kent noise.

Finally, we illustrate performance for a straight-line trajectory (Case III in Sec. III-C). The initial point is chosen uniformly randomly and velocity $\mathbf{v} = [0.1 \ -0.1 \ 0.2]^T$. The bearing observations y are sampled from a Kent distribution with mean equal to the true bearing. The range observations, s , are corrupted with additive white Gaussian noise to give an SNR of 30 dB. Figure 6 shows tracking performance averaged over 1000 realizations for various levels of noise in bearing and range observations.




Figure 2: Tracking performance with a constant rotation (Case I) averaged over 1000 realizations. The observation rate is 20 samples per second.

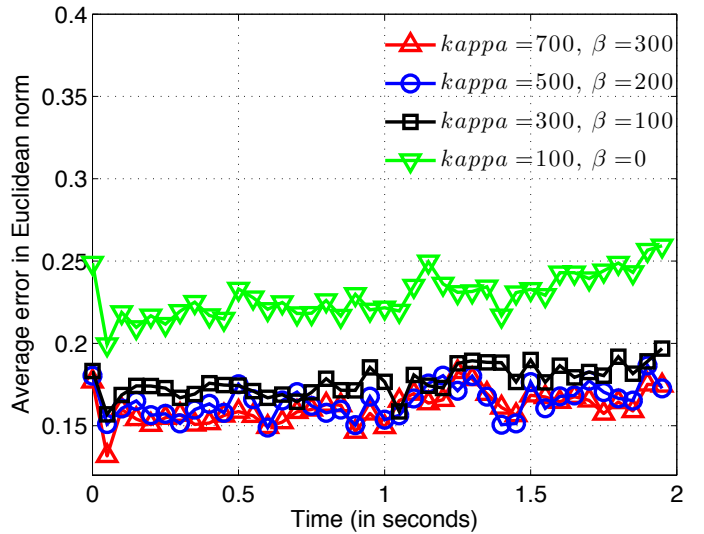


Figure 3: Tracking performance with a speeding target (changing rotation, Case II) averaged over 200 realizations. The observation rate is 20 samples per second.

Figures 1, 4, 5 show illustrations of motion, measurements, and our estimates for the three cases discussed in Section III.

VII. CONCLUSIONS

This paper presented an algorithm for tracking that alternates between learning rotations and translation in a sequential fashion. The algorithm leverages recent advances in learning rotations and separation of the state-space by decomposing the states in \mathbb{R}^k into states on $S^{k-1} \times \mathbb{R}$. This decomposition also respects the geometry of the observations in the spherical coordinates, namely bearing and range. This allows development

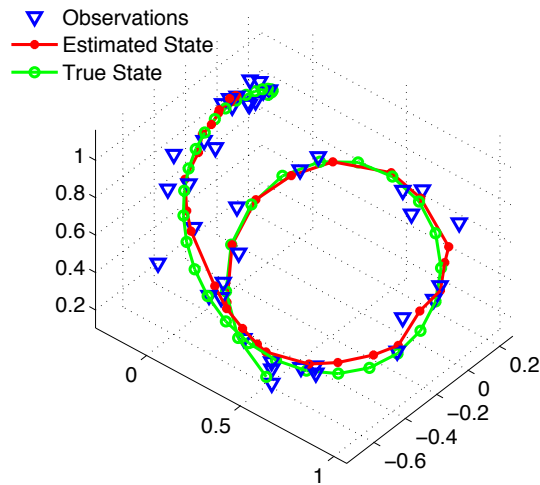


Figure 4: Example Simulation of Case II: object moving on sphere of radius 1 units from observer located at origin. The true motion shown in green involves changing rotation but no translation. The noisy bearing measurements (with noise simulated as Kent noise) are shown in blue and the track estimated from these noisy measurements by proposed updates is shown in red.

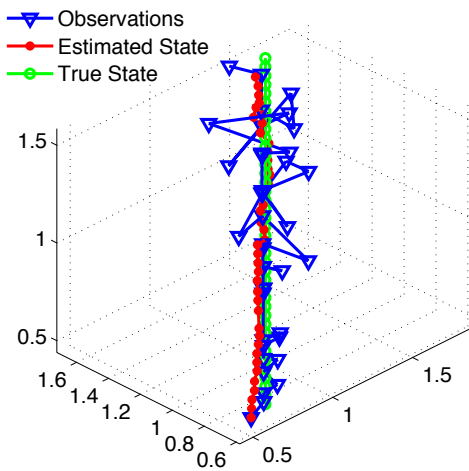


Figure 5: Example Simulation of Case III: object moving in a straight line. True motion shown in green, noisy measurements in blue and estimated track in red.

of unbiased estimators in the respective domains. Preliminary experimental results illustrate that the proposed method can work well. Future work will include comparisons to other

trackers and experiments with real data.




Figure 6: Tracking performance with changing rotation and translation (Case III) averaged over 1000 realizations. The observation rate is 20 samples per second.

REFERENCES

- [1] Thomas Kailath, "An innovation approach to least-squares estimation part I: Linear filtering in additive white noise," *IEEE Transactions on Automatic Control*, vol. 13, no. 6, pp. 646–655, 1968.
- [2] G. Welch and G. Bishop, "An introduction to the Kalman filter," Technical report TR 95-041, University of North Carolina, Department of Computer Science, 1995.
- [3] S. J. Julier and J.K. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *Intl. Symp. Aerospace/Defense Sensing, Simul. and Controls 3*, 1997.
- [4] Zhansheng Duan and X. Rong Li, "Sequential nonlinear tracking filter with range-rate measurements in spherical coordinates," in *Proc. of the 7th Intl. Conf. Information Fusion*, 2004, pp. 599–605.
- [5] John T. Kent, "The Fisher-Bingham distribution on the sphere," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 44, no. 1, pp. 71–80, 1982.
- [6] Adam Smith and Manfred Warmuth, "Learning rotations," in *Conference on Learning Theory*, Jun 2008.
- [7] Raman Arora, "On learning rotations," in *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- [8] R. Arora and W. A. Sethares, "An efficient and stable algorithm for learning rotations," in *Proc. Intl. Conf. Pattern Recognition (ICPR)*, 2010.
- [9] S. Fiori and T. Tanaka, "An algorithm to compute averages on matrix lie groups," *Signal Processing, IEEE Transactions on*, vol. 57, no. 12, pp. 4734–4743, dec. 2009.
- [10] Raman Arora, *Group theoretical methods in signal processing: learning similarities, transformations and invariants*, Ph.D. thesis, Univ. of Wisconsin-Madison, 2009.