A Particle Filter Localisation System for Indoor Track Cycling Using an Intrinsic Coordinate Model

Jiaming Liang and Simon Godsill
Department of Engineering
University of Cambridge
Cambridge, UK
Email: {jl809, sjg}@eng.cam.ac.uk

Abstract—In this paper we address the challenging task of tracking a fast-moving bicycle, in the indoor Velodrome environment, using inertial sensors and infrequent position measurements. Since the inertial sensors are physically in the intrinsic frame of the bike, we adopt an intrinsic frame dynamic model for the motion, based on curvilinear dynamical models for manoeuvring objects. We show that the combination of inertial measurements with the intrinsic dynamic model leads to linear equations, which may be incorporated effectively into particle filtering schemes. Position measurements are provided through timing measurements on the track from a camera-based system and these are fused with the inertial measurements using a particle filter weighting scheme. The proposed methods are evaluated on synthesised cycling datasets based on real motion trajectories, showing their potential accuracy, and then real data experiments are reported.

Index Terms—Object localisation, Inertial measurements, Sequential Monte Carlo, Track cycling, Intrinsic coordinate model, Curvilinear model

I. INTRODUCTION

In tracking problems ranging from pedestrian positioning to track cycling, the aim is to estimate the kinematic state of a target (position, orientation, etc.) over time according to a set of noisy measurements. These measurements may be provided by sensors such as inertial measurement units (IMUs), magnetometers, odometers, wireless signal receivers or GPS. Significant efforts have been devoted to the design and building of indoor positioning systems that are capable of providing pedestrians or vehicles with accurate location information in environments where GPS signals are unavailable [1]–[3]. Among these systems, infrastructures including ultra-wideband, infrared, RFID and Bluetooth are usually deployed to reduce position errors resulting from low-cost sensors. To fuse information from multiple information sources and build an integrated navigation system, Kalman filter-based methods and particle filters are possible candidate methods. Compared with the classical Kalman filter and its variants, particle filters have shown advantages in dealing with non-linear and non-Gaussian filtering problems [4], [5]. Here, our work aims to design novel particle filter-based methods to tackle on-line tracking for a fast moving object (in this case a bicycle) in an indoor environment using inertial sensors and occasional position measurements provided by timing lines on the Velodrome track.

We here adopt a dynamic model-based approach to the fusion of inertial measurements with position (timing) measurements, in an on-line localisation/tracking scenario. The idea is that an accurate modelling of the inertial sensors and their bias provides approximate information on the path, while constraints of track geometry and timing line measurements can be fused with the inertial data to achieve a refined tracking of position and orientation as it evolves over time. Rather than adopting standard motion models such as the white noise-driven constant-acceleration and constant-velocity models [6], [7], which may not accurately model object manoeuvres, we use a curvilinear dynamical model defined in the intrinsic coordinte frame of the object (see e.g. the excellent survey in [8], which points out that the standard random process based models may not accurately model the physical manoevres of an object; see also [9] which suggests that the appropriate use of target kinematic modelling can assist in an inertial tracking system.)

The class of curvilinear models that we propose is based on an intrinsic coordinate frame, as in the particle filter methods of [10], [11]; see also the similar model in [12], in which an interacting multiple model tracker is implemented for an intrinsic coordinate frame, and [13] which describes the dynamics of an aircraft using Frenet-Serret formulae, which can be regarded as a variation to the intrinsic coordinate model. While being powerful in terms of modelling target manoeuvres accurately, the intrinsic model turns out also to be very well adapted to the case of inertial measurements, since it leads to a set of linear, Gaussian equations for turn rate and ground speed measurements. The novelty of our work here lies in developing the analytic structure for the inertial measurements, and in blending this structure with a particle filter for fusing with the timing measurements. In particular we are able to make very effective hybrid proposals to the particle filter based upon linear update equations for speed and angle rate measurements, which achieves significantly better performance than an equivalent bootstrap SIR filter [4]. We also show a backward smoothing-based scheme that draws samples from the joint distribution of states in the recent and more distant past. The approaches proposed are then evaluated on synthetic data as well as real world data.

The remainder of the paper is organised as follows. In the next Section, we present the motion model using an intrinsic
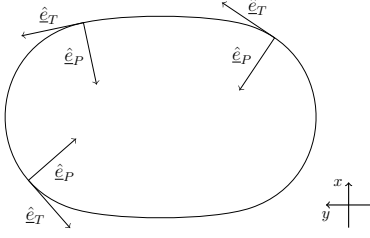
Fig. 1. 2D curvilinear motion in an intrinsic coordinate system. $\hat{\underline{e}}_T$ and $\hat{\underline{e}}_P$ depict tangential and perpendicular unit vectors respectively.

coordinate model. We then present the measurement models to link the sensor data to the dynamical model. Following this, two novel particle filter weight update schemes are described in Section III. Evaluation results on both synthetic and real data can be found in Section IV and conclusions are drawn in Section V.

II. DYNAMIC AND MEASUREMENT MODELS

In this section we describe the dynamical model, defined in terms of an intrinsic coordinate system, to express the motion of indoor Velodrome riders, and the corresponding measurements processes for the inertial, speed and position data.

A. Dynamical Model in an Intrinsic Coordinate System

The model here is based on [10], [11] and [14] and similar to curvilinear models in [8] and [12]. An applied force acting upon an object can be decomposed into a tangential component T_T and a component T_P perpendicular to the tangential vector. Let l denote the arc length along the curve on which a particle of mass m travels, and ψ the heading angle relative to a reference axis. The tangential and perpendicular equations of motion are then given by:

$$T_T = \lambda \frac{dl}{dt} + m \frac{d^2 l}{dt^2}, \tag{1}$$

$$T_P = m \frac{dl}{dt} \frac{d\psi}{dt}, \tag{2}$$

where $\lambda \frac{dl}{dt}$ is a resistance term. Assume that T_T and T_P are piecewise constant between time τ and $\tau + \Delta\tau$. A 2-dimensional motion may be described by the evolution of the tangential and perpendicular unit vectors for the object over time, see for example Fig. 1, which illustrates their motion around a velodrome track.

(1) may be solved directly as in [10], [14] to obtain an expression for speed $s_t = dl/dt$:

$$s_{\tau+\Delta\tau} = s_\tau e^{-\frac{\Delta\tau\lambda}{m}} + \frac{T_T}{\lambda}\left(1 - e^{-\frac{\Delta\tau\lambda}{m}}\right), \tag{3}$$

and l_t can then be routinely obtained by further integrating this expression [10], [14]. Now, using this result (2) can be solved for the change in heading angle and unit vectors:

$$\begin{bmatrix} \hat{\underline{e}}_{T,\tau+\Delta\tau} \\ \hat{\underline{e}}_{P,\tau+\Delta\tau} \end{bmatrix} = \begin{bmatrix} \cos(\Delta\psi) & \sin(\Delta\psi) \\ -\sin(\Delta\psi) & \cos(\Delta\psi) \end{bmatrix} \begin{bmatrix} \hat{\underline{e}}_{T,\tau} \\ \hat{\underline{e}}_{P,\tau} \end{bmatrix}, \tag{4}$$

$$\Delta\psi = \frac{T_P}{T_T}\left(\Delta\tau\lambda/m - \log\left|\frac{s(\tau)}{s(\tau+\Delta\tau)}\right|\right). \tag{5}$$

Finally, the change in Cartesian position can be computed in closed form only for the case $\lambda = 0$ (no resistance to motion) [11], so in general the position $z_t = [x_t, y_t]^T$ can be obtained for example by a fine-grained Euler approximation:

$$z_{\tau+\Delta\tau} \approx z_\tau + \Delta\tau \cdot s(\tau) \cdot \hat{\underline{e}}_{T,\tau}. \tag{6}$$

B. Speed Measurement Model

If the tangential thrust is assumed to be drawn from a simple Gaussian distribution, say $T_T \sim \mathcal{N}(\mu_T, \sigma_T^2)$, (3) becomes a linear Gaussian state space model and the transition density of speed is as follows:

$$f(s_{\tau+\Delta\tau}|s_\tau) = \mathcal{N}\Big(s_{\tau+\Delta\tau}|e^{-\frac{\Delta\tau\lambda}{m}}s_\tau + \frac{\mu_T}{\lambda}\left(1 - e^{-\frac{\Delta\tau\lambda}{m}}\right), $$
$$\frac{\sigma_T^2}{\lambda^2}\left(1 - e^{-\frac{\Delta\tau\lambda}{m}}\right)^2\Big). \tag{7}$$

Modeling the speed in this particular way allows the use of linear Kalman-style updating when direct speed measurements are available:

$$\hat{s}_\tau = s_\tau + n_{\hat{s},\tau}, \tag{8}$$

where the measurement noise $n_{\hat{s},\tau}$ is additive, white Gaussian noise of the form $n_{s,\tau} \sim \mathcal{N}(0, \sigma_{\hat{s}}^2)$.

C. Heading Measurement Model

Conditioned on the speed of the object, (2) can also be rearranged to give the dynamic equation of heading rate

$$\dot{\psi}_t = \frac{T_P}{m s_t}, \quad \tau < t < \tau + \Delta\tau, \quad s_t \neq 0, \tag{9}$$

This is again a linear Gaussian model conditioning on s_τ, under the assumption that $T_P \sim \mathcal{N}(0, \sigma_P^2)$:

$$f(\dot{\psi}_\tau|s_\tau) = \mathcal{N}\Big(\dot{\psi}_\tau|0, \frac{\sigma_P^2}{m^2 s_\tau^2}\Big). \tag{10}$$

Now, given direct noisy measurements of the turn rate of the object in the 2D plane, $\hat{\dot{\psi}}_\tau$, which include a random additive bias term b_t, with Gaussian random walk transition density:

$$b_{\tau+\Delta_\tau} = b_\tau + \Delta\tau \cdot n_{b,\tau}, \ n_{b,\tau} \sim \mathcal{N}(0, \sigma_b^2), \tag{11}$$

the observation model for turn rate is given by:

$$\hat{\dot{\psi}}_\tau = \dot{\psi}_\tau + b_\tau + n_{\hat{\dot{\psi}},\tau}, \ n_{\hat{\dot{\psi}},\tau} \sim \mathcal{N}\left(0, \sigma_{\hat{\dot{\psi}}}^2\right). \tag{12}$$

The orientation of the IMU relative to the bike is shown in Fig. 2. When cycling around a track, direct measurements of heading rate may be unavailable. This is because the geometry of the track and the bike sway prevent the vertical axis of

gyroscope (i.e. $Z+$ in Fig. 2) from being aligned with the normal vector of the plane of motion. This suggests that we could obtain angular velocity on the plane of motion by rotating the gyroscope's frame onto the intrinsic frame, a development that will be reported in future work. In this paper however we treat the problem as a purely 2D tracking model. As a result, angular speed from the gyroscope's Z-axis is treated as a direct measurement of the rate of heading. The bias term b_t in our observation model will then include a component both from the modelling error (assumed to be small if the cycle is close to vertical) and the true instrumentation bias of the device.

D. Position Measurement Model

As discussed in the previous section, a major issue with the consumer-grade inertial sensors such as gyroscopes and accelerometers is that they usually suffer from bias. To deal with drift resulting from an inertial navigation system, other positioning systems such as GPS may be integrated. However, unlike outdoor sports where it is possible to take advantage of global navigation systems to help localise an object, we need additional positioning systems for indoor track cycling. Here, two additional systems are available for providing position information:

Timings of Crossing Lines and Lateral Position: There are 9 lines (0m, 25m, 50m, 100m, 115m, 125m, 150m, 200m and 240m) across the track, see Fig. 3, and the crossing times for these lines are measured. Such measurements are treated as a Gaussian measurement of the true position, $\hat{z}_\tau$ at the measured crossing time τ, truncated to lie within the track width:

$$f(\hat{z}_\tau | z_\tau) \propto \mathcal{N}(\hat{z}_\tau | z_\tau, \Sigma_z) \tag{13}$$

where $\hat{z}_\tau$ and Σ_z are respectively centred on and aligned with each timing line. When accurate position information is not available, Σ_z would be set very wide along the timing line. Fig. 3 shows the positions of timing lines along the track and the corresponding Gaussian distributions for two of the timing lines. For certain timing lines along the straights
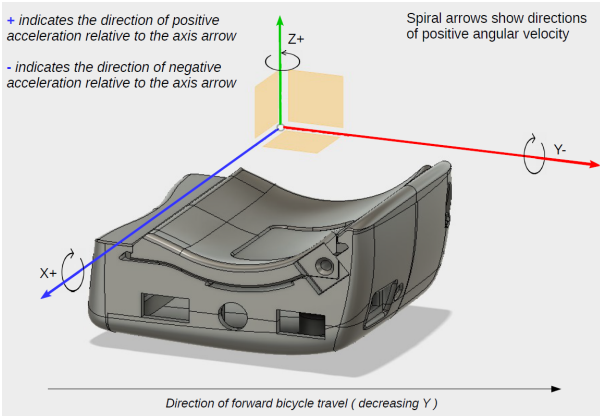

Fig. 2. Bike-mounted IMU orientation: $Z+$ is alligned to the surface normal when riding on flat ground
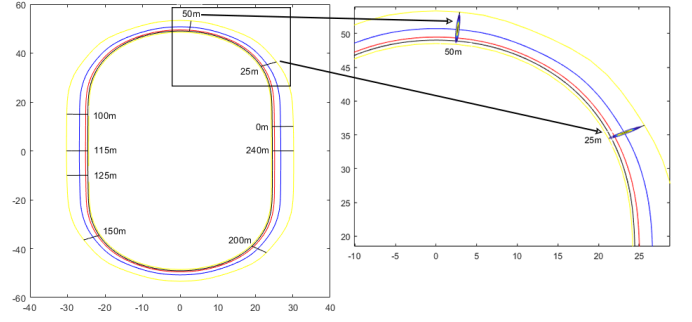

Fig. 3. Left: Top view of the track with timing lines plotted in black. Yellow contours represent the edge of velodrome track. Right: Gaussian distributions on 25m and 50m timing lines

(i.e. 0m, 100m, 115m, 125m and 240m) a camera system measuring lateral distance to the track's inner black line has been deployed around the velodrome, see [15]. In these cases the same truncated Gaussian is employed, but with a narrower covariance function centred on the measured lateral position.

III. STATE ESTIMATION USING SEQUENTIAL MONTE CARLO METHODS

Since we are dealing with a tracking problem with non-linear dynamical model, there is no closed-form analytic solution for the whole problem. Therefore, we consider using sequential Monte Carlo (SMC) methods, or particle filters [4], [16], [17], to sequentially estimate the target states under a Bayesian framework.

In this section, two approaches for improving performance are presented: the first cascades locally optimal proposal kernels together, seeking to reduce the mismatch between prior predictive distribution and the posterior distribution conditioned on new measurements, and the second provides a section-wise Kalman-smoothing procedure for the particle filter. Note that t and $t-1$ are used to replace $\tau + \Delta\tau$ and τ respectively in the following derivation. Also $\{x_{0:t}^{(i)}\}_{1 \le i \le N}$ is used to denote samples drawn from a importance distribution $q(x_{0:t}|y_{0:t})$.

A. Optimal Kernel Scheme

Owing to the particular design of our models, we can sample states from efficiently constructed locally optimal kernels. Specifically, for speed we suggest an importance distribution $q(\cdot)$ in the following form, using (7), (8):

$$q(s_t | \hat{s}_t, s_{t-1}^{(i)}) = p(s_t | \hat{s}_t, s_{t-1}^{(i)}) \propto p(\hat{s}_t | s_t) p(s_t | s_{t-1}^{(i)})$$
$$= \mathcal{N}(s_t | m_s, \sigma_s^2), \tag{14}$$

with

$$m_s = a s_{t-1}^{(i)} + b + \sigma_v^2(\hat{s}_t - a s_{t-1}^{(i)} - b)/(\sigma_v^2 + \sigma_{\hat{s}}^2),$$
$$\sigma_s^2 = \sigma_v^2 - \sigma_v^4/(\sigma_v^2 + \sigma_{\hat{s}}^2),$$
$$a = e^{-\Delta\tau\lambda/m} s_\tau, \ b = \mu_T(1 - e^{-\Delta\tau\lambda/m})/\lambda,$$
$$\sigma_v^2 = \sigma_T^2(1 - e^{-\Delta\tau\lambda/m})^2/\lambda^2.$$

This distribution is the *optimal kernel* for the speed in the sense of [17] conditioned solely upon the speed observation. Similarly, we suggest an importance distribution for $\boldsymbol{\alpha}_t = [\dot{\psi}_t, b_{\dot{\psi},t}]^T$ as follows:

$$
\begin{aligned}
q(\boldsymbol{\alpha}_t|\hat{\dot{\psi}}_t, \boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)}) &= p(\boldsymbol{\alpha}_t|\hat{\dot{\psi}}_t, \boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)}) \\
&\propto p(\hat{\dot{\psi}}_t|\boldsymbol{\alpha}_t, \boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)})p(\boldsymbol{\alpha}_t|\boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)}) \\
&= \mathcal{N}(\hat{\dot{\psi}}_t|B\dot{\psi}_t, \sigma_{\dot{\psi}}^2)\mathcal{N}(\boldsymbol{\alpha}_t|A\boldsymbol{\alpha}_{t-1}^{(i)}, C) \\
&= \mathcal{N}(\boldsymbol{\alpha}_t|m_{\boldsymbol{\alpha}}, P_{\boldsymbol{\alpha}}),
\end{aligned} \quad (15)
$$

where

$$
m_{\boldsymbol{\alpha}} = A\boldsymbol{\alpha}_{t-1}^{(i)} + CB^T(BCB^T + \sigma_{\dot{\psi}}^2)^{-1}(\hat{\dot{\psi}}_t - BA\boldsymbol{\alpha}_{t-1}^{(i)}),
$$

$$
P_{\boldsymbol{\alpha}} = C - CB(CB^T(BCB^T + \sigma_{\dot{\psi}}^2)^{-1}BC,
$$

$$
A = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \ B = [1 \ 1],
$$

$$
C = \begin{bmatrix} \sigma_P^2/m^2 s_{\tau+\Delta\tau}^2 & 0 \\ 0 & \Delta\tau^2\sigma_b^2 \end{bmatrix}.
$$

Once again, this is a locally *optimal kernel* conditioned on the speed $s_t^{(i)}$ and the heading rate measurement.

Finally, sequential importance sampling (SIS) can be applied to obtain samples from $\pi(\boldsymbol{\alpha}_{0:t}, s_{0:t}|\hat{\dot{\psi}}_{0:t}, \hat{s}_{0:t}, \hat{z}_{0:t})$. The update equation for the unnormalised importance weight $\tilde{w}_t^{(i)}$ is given as:

$$
\begin{aligned}
\tilde{w}_t^{(i)} &= \frac{\pi(\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}|\hat{\dot{\psi}}_{0:t}, \hat{s}_{0:t}, \hat{z}_{0:t})}{q(\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}|\hat{\dot{\psi}}_{0:t}, \hat{s}_{0:t})} \\
&= \tilde{w}_{t-1}^{(i)} \times \frac{p(\hat{\dot{\psi}}_t, \hat{s}_t|\boldsymbol{\alpha}_t^{(i)}, s_t^{(i)})p(\boldsymbol{\alpha}_t^{(i)}, s_t^{(i)}|\boldsymbol{\alpha}_{t-1}^{(i)}, s_{t-1}^{(i)})}{p(s_t^{(i)}|\hat{s}_t, s_{t-1}^{(i)})p(\boldsymbol{\alpha}_t^{(i)}|\hat{\dot{\psi}}_t, \boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)})} \\
&\quad \times \frac{p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)})}{p(\hat{\dot{\psi}}_t, \hat{s}_t, \hat{z}_t|\hat{\dot{\psi}}_{0:t-1}, \hat{s}_{0:t-1}, \hat{z}_{0:t-1})} \\
&\propto \tilde{w}_{t-1}^{(i)} \times p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)})p(\hat{s}_t|s_{t-1}^{(i)})p(\hat{\dot{\psi}}_t|\boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)}) \\
&\quad \times \frac{p(\hat{\dot{\psi}}_t|\boldsymbol{\alpha}_t^{(i)})p(\hat{s}_t|s_t^{(i)})p(\boldsymbol{\alpha}_t^{(i)}|\boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)})p(s_t^{(i)}|s_{t-1}^{(i)})}{p(s_t^{(i)}|s_{t-1}^{(i)})p(\hat{s}_t|s_t^{(i)})p(\boldsymbol{\alpha}_t^{(i)}|\boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)})p(\hat{\dot{\psi}}_t|\boldsymbol{\alpha}_t^{(i)})} \\
&= \tilde{w}_{t-1}^{(i)} \times p(\hat{s}_t|s_{t-1}^{(i)})p(\hat{\dot{\psi}}_t|\boldsymbol{\alpha}_{t-1}^{(i)}, s_t^{(i)})p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}).
\end{aligned} \quad (16)
$$

Note that in the velodrome tracking scenario where timing measurements are not available at every timestamp we can compute the likelihood $p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)})$ only when the object passes one of the timing lines - at all other times the term $p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)})$ is omitted from the weight update. When timing measurements are present, the likelihood is computed as $p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}) = f(\hat{z}_t|z_t^{(i)})$ as in (13) and $z_t^{(i)}$ is the Cartesian location of the i-th particle, obtained using (6).

B. Section-wise Backward Smoothing Scheme

Here we propose a scheme that performs section-wise backward smoothing, aimed at improving the retrospective performance of the scheme and in enhancing the proposed state sequences. A natural partitioning of the data is into the 9 sections defined by the timing lines, although the scheme is operable over any interval desired. The smoothing is based on the observation that the speed states are jointly Gaussian conditioned on the speed measurements, i.e. $p(s_{0:T}|\hat{s}_{0:T})$ is a $T + 1$-variate Gaussian that may be sampled using standard forward-filtering backward sampling (FFBS) methods, see [18], [19]. This arises because of the linear Gaussian state-space structure of (7) and (8). Similarly, the linear Gaussian structure of (10) and (11) means that the heading rate and bias have a jointly Gaussian distribution conditioned on sampled speed and heading measurements, $p(\boldsymbol{\alpha}_{0:t}|\hat{\dot{\psi}}_{0:t}, s_{0:t}^{(i)})$, which may also be drawn jointly using FFBS. Recall that the FFBS scheme applied to speed factorises the joint smoothing distribution of speed, $p(s_{0:T}|\hat{s}_{0:T})$, as follows

$$
\begin{aligned}
p(s_{0:T}|\hat{s}_{0:T}) &= p(s_T|\hat{s}_{0:T})\prod_{t=0}^{T-1} p(s_t|s_{t+1:T}, \hat{s}_{0:T}) \\
&= p(s_T|\hat{s}_{0:T})\prod_{t=0}^{T-1} p(s_t|s_{t+1}, \hat{s}_{0:t}),
\end{aligned} \quad (17)
$$

and is easily adapted to a sub-section of the state sequence from t_1 to t_2.

The proposed scheme then draws a joint sequence of speeds from the the local optimal distribution $q(s_{t_1+1:t_2}|s_{t_1}) = p(s_{t_1+1:t_2}|\hat{s}_{t_1+1:t_2}, s_{t_1})$ using FFBS and then proposes a sequence of $\boldsymbol{\alpha}_t$ states conditionally upon that proposed sequence of speeds according to $q(\boldsymbol{\alpha}_{t_1+1:t_2}|\boldsymbol{\alpha}_{t_1}) = p(\boldsymbol{\alpha}_{t_1+1:t_2}|s_{t_1+1:t_2}, \hat{\dot{\psi}}_{t_1+1:t_2}, \boldsymbol{\alpha}_{t_1})$. The details of the FFBS scheme for each of these draws can be found in Appendix.

This proposal can be considered as a section-wise batch version of the locally optimal proposal from the previous section, and the weights are obtained in a similar fashion as:

$$
\begin{aligned}
\tilde{w}_t^{(i)} &= \frac{\pi(\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}|\hat{\dot{\psi}}_{0:t}, \hat{s}_{0:t}, \hat{z}_{0:t})}{q(\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}|\hat{\dot{\psi}}_{0:t}, \hat{s}_{0:t})} \\
&\propto \tilde{w}_{t_1}^{(i)} \frac{p(s_{t_1+1:t}^{(i)}|\hat{s}_{t_1+1:t}, s_{t_1}^{(i)})p(\boldsymbol{\alpha}_{t_1+1:t}^{(i)}|\hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)})}{p(s_{t_1+1:t}^{(i)}|\hat{s}_{t_1+1:t}, s_{t_1}^{(i)})p(\boldsymbol{\alpha}_{t_1+1:t}^{(i)}|\hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)})} \\
&\quad \times p(\hat{\dot{\psi}}_{t_1+1:t}|s_{t_1+1:t}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)})p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}) \\
&= \tilde{w}_{t_1}^{(i)}p(\hat{\dot{\psi}}_{t_1+1:t}|s_{t_1+1:t}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)})p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)}),
\end{aligned} \quad (18)
$$

given $t_1 < t \leq t_2$. Note that t_1 and t_2 are usually (but not necessarily) the timings at two consecutive timing lines while $\tilde{w}_{t_1}^{(i)}$ is the weight of particle i computed at time t_1 (i.e. at the end of the previous section). Again we only compute $p(\hat{z}_t|\boldsymbol{\alpha}_{0:t}^{(i)}, s_{0:t}^{(i)})$ when arriving at a timing line, typically at the end of the segment $t = t_2$. Also, the second term in (18) can be computed as

$$
\begin{aligned}
p(\hat{\dot{\psi}}_{t_1+1:t}|s_{t_1+1:t}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)}) &= p(\hat{\dot{\psi}}_{t_1+1}|s_{t_1+1}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)}) \\
&\times \prod_{k=t_1+2}^{t} p(\hat{\dot{\psi}}_k|\hat{\dot{\psi}}_{t_1+1:k-1}, s_{t_1+1:k}^{(i)}, \boldsymbol{\alpha}_{t_1}^{(i)}),
\end{aligned} \quad (19)
$$

where

$$p(\hat{\psi}_t|\hat{\psi}_{t_1+1:t-1}, s^{(i)}_{t_1+1:t}, \boldsymbol{\alpha}^{(i)}_{t_1}) \qquad (20)$$

$$= \int p(\hat{\psi}_t|\boldsymbol{\alpha}_t)p(\boldsymbol{\alpha}_t|\hat{\psi}_{t_1+1:t-1}, s^{(i)}_{t_1+1:t}, \boldsymbol{\alpha}^{(i)}_{t_1})d\boldsymbol{\alpha}_t$$

$$= \int \mathcal{N}(\hat{\psi}_t|B\boldsymbol{\alpha}_t, \sigma^2_{\hat{\psi}})p(\boldsymbol{\alpha}_t|\mu^{\boldsymbol{\alpha}}_{t|t-1}, P^{\boldsymbol{\alpha}}_{t|t-1})d\boldsymbol{\alpha}_t$$

$$= \mathcal{N}(\hat{\psi}_t|B\mu^{\boldsymbol{\alpha}}_{t|t-1}, BP^{\boldsymbol{\alpha}}_{t|t-1}B^T + \sigma^2_{\hat{\psi}}), \qquad (21)$$

with the predictive distribution $p(\boldsymbol{\alpha}_t|\mu^{\boldsymbol{\alpha}}_{t|t-1}, P^{\boldsymbol{\alpha}}_{t|t-1})$ provided by a Kalman filter running on the i-th particle. Based on (18), SIS can be implemented again to draw samples from the target distribution.

C. On-track Detection

Apart from aforementioned observations, we also possess point cloud data acquired by laser scanning. This type of mapping information allows us to determine whether or not a particle stays physically within the track width. We can then assign zero weight to an out-of-track particle and resample particles when, for example, the effective sample size (ESS) is lower than a pre-defined threshold. Currently this detection mechanism is implemented only in the optimal kernel scheme in which it is easy to detect whether a particle will go off-track at the next time epoch. For the sectionwise smoothing scheme, we will not know whether a particle is moving off-track until the sectionwise backward smoothing step is completed. This more elaborate constraint enforcement is left as a topic of future exploration.

IV. EXPERIMENTAL RESULTS

As the ground truth for the real data is not available, we synthesise several different datasets, each of which mimics a 2D cycling trajectory on the same Velodrome track. By doing this, we could test and validate the tracking performance of our methods. After this, an experiment is conducted by applying the proposed methods to a real track cycling dataset in which the cycling trajectory is pre-defined. All tests would be implemented in MATLAB without code optimisation.

A. Results on Synthetic Data

Our idea of synthesising track cycling data is to run particle filters on the measurements acquired during the time when a cyclist cycled around a track for several laps. Then we adopt the output of the filters as our ground truth for further use. The dynamic and the measurements models are those defined in Section II and the parameters we use to generate the synthetic data are given in Table I. Also note that the sensor measurements are logged every 0.1 seconds in order to be in line with the actual scenario. In total, there are 5 datasets, each corresponding to a 2- or 3-lap on-track trajectory.

We now present performance comparison based on RMSE values between different schemes. It is worth noting that in practice we only have timing information rather than the relative distance between the cyclist and the black measurement line at corner timing lines. However, we start with assuming

TABLE I
PARAMETERS USED TO GENERATE SYNTHETIC DATA

λ	m	μ_T	μ_P	σ_T	σ_P	σ_b	$\sigma_{\hat{s}}$	$\sigma_{\hat{\psi}}$
1	100	0	0	50	3000	0.5°	0.5	18°

TABLE II
OVERALL TESTING RESULTS FOR 5 DATASETS (10 RUNS FOR EACH SCHEME). NUMBERS GIVEN ARE MEANS (AND STANDARD DEVIATIONS) OF POSITION RMSES

		filtering RMSEs	fix-lag smth. RMSEs	fix-int. smth. RMSEs
	#particles	position [m]	position [m]	position [m]
OPT	500	1.691 (0.236)	0.958 (0.216)	0.773 (0.165)
	1000	1.550 (0.243)	0.817 (0.177)	0.613 (0.113)
	2000	1.546 (0.134)	0.776 (0.123)	**0.526** (0.055)
FFBS	100	4.606 (0.727)	2.475 (0.485)	2.021 (0.523)
	150	4.299 (0.747)	2.199 (0.686)	1.523 (0.767)
	300	3.547 (0.350)	1.451 (0.198)	**0.811** (0.097)
SIR	2000	5.481 (0.545)	4.237 (0.564)	**3.993** (0.568)

such measurements are available in our first experiment whose results are showed in Table II. Later we will present tracking results with only corner timing measurements to evaluate the necessity of having corner position measurements. Also, the parameters have been adjusted manually to get the best possible results in all schemes.

Specifically, the schemes are tested 10 times on each of the synthesised datasets. The overall error performance in position is given in Table II. Three measures of error are compared: the first one uses the filtering state estimate, the second one uses the fix-lag smoothing state estimate ($lag = 15$) and the last one the fix-interval smoothing state estimate. Different amount of particles are used in the tests in order to compare performance between schemes given similar runtime. For instance, the computation time of running 2000 particles in the optimal kernel scheme (OPT) would be similar to that of running 300 particles in the backward smoothing-based scheme (FFBS). A bootstrap SIR filter (SIR) with track geometry constraint is also run on the datasets to serve as a benchmark. As shown in the table, the results for the OPT outperforms both of the other two schemes. One of the reasons that the FFBS scheme is less effective than the OPT schemes is its lack of an on-track detection mechanism. Particles with non-zero weights operating outside the track between timing lines introduce negative contributions to the state estimation. However, the FFBS scheme still exceeds the bootstrap filter. The bootstrap filter has the worst performance as it uses state transition density as its kernel. Drawing samples from the Gaussian prior results in more diverse heading rates and thus worse tracking performance. It is common that in the bootstrap filter many particles fail to get to timing lines in

#particle	OPT			FFBS			SIR
	500	1000	2000	100	150	300	2000
	position [m]			position [m]			position [m]
1	0.601 (0.161)	0.551 (0.134)	0.470 (0.098)	1.050 (0.676)	0.999 (0.481)	0.679 (0.149)	2.158 (0.638)
2	0.872 (0.170)	0.662 (0.104)	0.570 (0.068)	2.513 (1.852)	1.022 (0.502)	0.759 (0.177)	2.543 (0.945)
3	0.651 (0.139)	0.560 (0.069)	0.525 (0.052)	2.3081 (1.125)	1.881 (1.611)	0.856 (0.109)	4.941 (1.678)
4	0.756 (0.215)	0.679 (0.332)	0.485 (0.071)	1.358 (0.806)	1.637 (1.505)	0.869 (0.298)	3.187 (0.611)
5	0.922 (0.411)	0.575 (0.093)	0.557 (0.155)	1.245 (0.682)	1.709 (1.753)	0.738 (0.222)	4.644 (0.876)

time, which further degrades the performance of the filter. In contrast, the proposed schemes perform better as the mismatch between prior predictive distribution and posterior distribution is reduced. Furthermore, results regarding different measures of error suggest that the fix-interval smoothing state estimation is more appropriate for the scenario of indoor tracking cycling given that the width of the track is about 7 meters.

More details on tracking results regarding individual dataset can be found in Table III. Here we only present the RMSEs obtained by the fix-interval smoothing state estimation. Again, the OPT scheme surpasses the others in terms of the individual position error performance. The results for the bootstrap filter are dominated by a number of cases with very large position errors. Moreover, the performance of the proposed schemes are getting better when the number of particles is increased. This implies that a tradeoff between performance and particle number has to be made when considering running the filters in an on-line sense. To find out different running time against different numbers of particles, please see Table IV. The numbers in entries denote the ratio of runtime over actual time. For the OPT scheme, the on-track detection that is based on an in-polygon test accounts for about 38 percent of the total runtime according to the MATLAB *profile*, which suggests a course of improvement in the future. As for the FFBS scheme, one major portion of its computation cost is caused by running N (number of particles) Kalman filters simultaneously.

We now investigate the situation where lateral position measurements are not available on corner timing lines. Constrained Gaussian distributions with very large covariance are deployed on these lines (i.e. 25m, 50m, 150m and 200m lines), which shows no preference for particles that are able to cross one of the lines in time. Tracking results given in Table V reveal that all filters suffer from reduction of performance when only timing information is available on the corner lines. Particularly, less frequent position measurements increase the chances of having particles getting out of track for the FFBS scheme, which sometimes leads to poor performance. On the other hand, the performance of the OPT scheme has shown that it is relatively robust even though less position measurements are provided. To give a better visualisation with respect to the tracking performance for each scheme, we show the time evo-

Methods	OPT			FFBS			SIR
#particles	500	1000	2000	100	150	300	2000
	ratios (runtime/actual time)						
1 (37.4s)	1.72	3.14	6.15	1.71	2.53	4.84	5.26
2 (29.4s)	1.53	3.04	6.00	1.55	2.25	4.38	5.23
3 (62.5s)	1.65	3.28	6.46	1.83	2.65	5.25	5.68
4 (41.3s)	1.52	2.97	5.96	1.73	2.58	5.08	5.15
5 (41.7s)	1.50	2.98	5.96	2.05	2.59	5.10	5.14

lution of the position errors and the corresponding empirical cumulative density functions regarding dataset 1 in Fig. 4. The blue and red vertical lines on the left denote the timings of crossing corner and straight timing lines respectively. It can be seen from the top left graph that the proposed schemes can keep the estimation errors at low levels when corner position measurements are provided. In this case, 85% of the position errors for the OPT scheme are below 0.616m while the same percent of the errors in the FFBS scheme lie below 0.728m. As for the case where only timings are provided around the corner, the OPT schemes still keeps 85% of the errors under 0.896m. And it becomes more clear that without the on-track constraint and the corner position measurements, the performance of the FFBS scheme drops.

B. Preliminary Results on Real Data

As mentioned, the lateral position measurements are not available in practice and there is no groundtruth. To cope with this situation, we firstly ask the cyclist to ride a bike along a pre-defined trajectory on the track. The path is defined as follows:

1) from 115m timing line to 35m line, the cyclist stays on the black measurement line.
2) at 35m line, the cyclist starts a vector from the black measurement and end up with around 220cm (slightly

TABLE V
WITH/WITHOUT CORNER POSITION MEASUREMENT (10 RUNS), PARTICLES
USED IN THREE SCHEMES ARE: 2000, 300 AND 2000 RESPECTIVELY.
NUMBERS GIVEN ARE MEANS (AND STANDARD DEVIATIONS) OF
POSITION RMSES

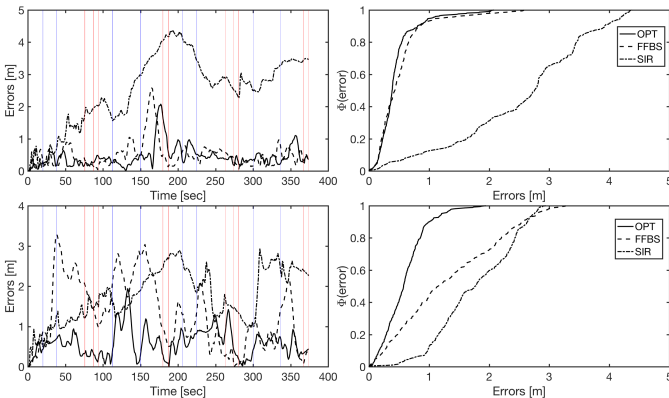|	with corner position?	RMSE position [m]
OPT	no	0.681 (0.071)
	yes	0.526 (0.055)
FFBS	no	1.771 (0.462)
	yes	0.811 (0.097)
SIR	no	3.992 (0.781)
	yes	3.993 (0.568)

Fig. 4. Tracking performance of different schemes on Dataset 1. Left: position errors along time, with(top left)/without(bottom left) corner position measurements. Blue vertical lines: timings at the corner timing lines; Red vertical lines: timings at the timing lines on the straights. Right: empirical CDFs of position errors, with(top right)/without(bottom right) corner position measurements.

below the blue line which is 250cm away from the black measurement line) at 50m timing line.

3) from 50m timing line on, the cyclist dips down and cycles around the red line.
4) from 110m line (i.e. between 100m and 115m timing line) to 115m timing line, the cyclist drops again from the red line to the black line
5) repeat for a few laps.

Lateral position measurements and timing measurements are synchronised with the on-bike sensor measurements. The sampling rate is set to be 10Hz for the on-bike sensors. On the corner timing lines, constrained Gaussian distributions, which are the same as the ones used in the synthetic datasets, are implemented. We then test our proposed methods on a collected cycling dataset.

As shown in Fig. 5, the OPT scheme gives a tracking result which matches the defined trajectory well. 95% confidence ellipses are occasionally plotted along the estimated trajectory. From the magnified area 2-3 (i.e. corresponding to the step 2
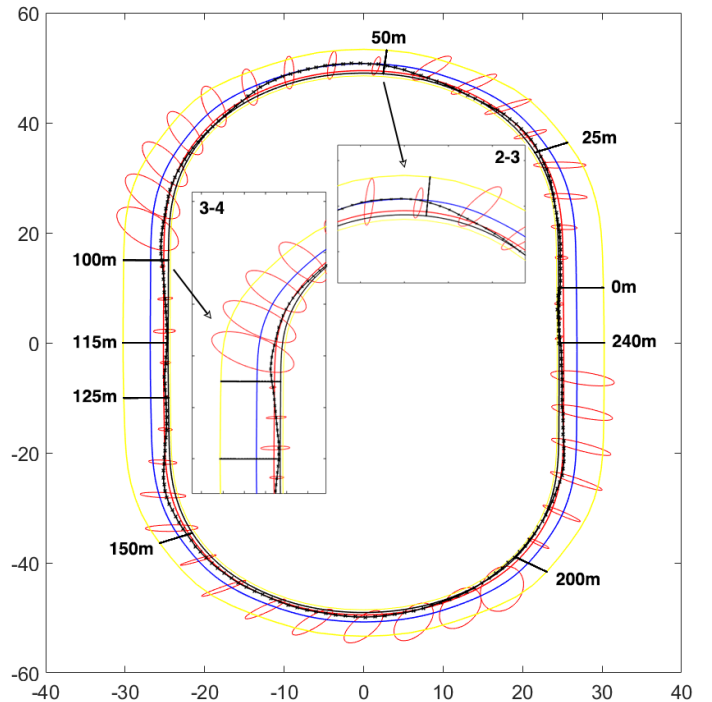

Fig. 5. An 1-lap trajectory estimated by the OPT scheme and 95% confidence ellipses along time.

and 3), it can be seen that the ellipse becomes smaller as just a subset of the particles get to the line in time. As the particles are moving between the 50m line and the 100m line, the uncertainty becomes gradually larger as the model relies only on the IMU and speed measurements. However, owing to the optimal kernels and the on-track constraint the ellipses can still cover the red line around which the rider is positioned. This guarantees that the tracking model can use the accurate lateral position measurements to get back to the correct path when the rider arrives at the 100m timing line. As to the FFBS scheme, although there are more uncertainty around the corner due to the unconstrained particle motion it can still follow the predefined path. This is because the correct positions are still within the error ellipses. The example trajectory obtained by the FFBS scheme is not plotted here as it is visually very similar to that of the OPT scheme.

V. CONCLUSION

In this work, we have proposed two novel schemes which allow an improved use of the model structure compared to simpler particle filtering schemes. Experimental results on synthetic cycling datasets have shown in particular that the tracking performance of our model is improved over a bootstrap SIR filter. Also, results on a real dataset show promise for future application of the methods.

It should be noted that there is scope for improving the models. For instance, a binormal vector, which is vertical to the plane formed by the tangential and perpendicular vector, could be added to obtain a 3D intrinsic coordinate system.

Given the fact that the cyclist would not physically leave the track surface under normal conditions, this binormal vector can be set to be equal to the surface normal, thus further constraining the 3D motion of a cyclist. Also, we could obtain more accurate information regarding the heading rate by tracking the tilt angle. Moreover, the tracking accuracy may be further improved by adopting existing sensors such as accelerometers - we note that similar equations for linear updating of accelerometer data apply to our intrinsic model, and these will be reported in future work.

APPENDIX

Here we give the equations used in the section-wise backward smoothing scheme. Supposing that $p(s_t|\hat{s}_{t_1+1:t}, s_{t_1}) = \mathcal{N}(s_t|\mu^s_{t|t}, \sigma^2_{t|t})$ is the output from a Kalman filter running on the i-th particle, we can compute the term $p(s_t|s_{t+1}, \hat{s}_{t_1+1:t}, s_{t_1})$ as below

$$p(s_t|s_{t+1}, \hat{s}_{t_1+1:t}, s_{t_1}) = \frac{p(s_{t+1}|s_t)p(s_t|\hat{s}_{t_1+1:t}, s_{t_1})}{p(s_{t+1}|\hat{s}_{t_1+1:t}, s_{t_1})}$$
$$= \mathcal{N}(s_t|\tilde{m}_s, \tilde{\sigma}^2_s), \qquad (22)$$

with

$$\tilde{m}_s = \mu^s_{t|t} + b + a\sigma^2_{t|t}(a^2\sigma^2_{t|t} + \sigma^2_v)^{-1}(s_{t+1} - a\mu^s_{t|t} - b),$$
$$\tilde{\sigma}^2_s = \sigma^2_{t|t} - a^2\sigma^4_{t|t}(a^2\sigma^2_{t|t} + \sigma^2_v)^{-1}.$$

Then we can firstly draw $s^{(i)}_{t_2}$ given the measurements $\hat{s}_{t_1+1:t_2}$ and then work backwards in time drawing state using the following equation:

$$p(s_{t_1+1:t_2}|\hat{s}_{t_1+1:t_2}, s_{t_1})$$
$$= p(s_{t_2}|\hat{s}_{t_1+1:t_2}, s_{t_1}) \prod_{t=t_1+1}^{t_2-1} p(s_t|s_{t+1}, \hat{s}_{t_1+1:t}, s_{t_1}). \quad (23)$$

It is possible to factorise the joint smoothing distribution for the heading rate and the bias in a similar fashion as:

$$p(\boldsymbol{\alpha}_{t_1+1:t_2}|\hat{\dot{\psi}}_{t_1+1:t_2}, s_{t_1+1:t_2}, \boldsymbol{\alpha}_{t_1})$$
$$= p(\boldsymbol{\alpha}_{t_2}|\hat{\dot{\psi}}_{t_1+1:t_2}, s_{t_1+1:t_2}, \boldsymbol{\alpha}_{t_1})$$
$$\times \prod_{t=t_1+1}^{t_2-1} p(\boldsymbol{\alpha}_t|\boldsymbol{\alpha}_{t+1}, \hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t}, \boldsymbol{\alpha}_{t_1}). \quad (24)$$

with the terms inside the product computed as

$$p(\boldsymbol{\alpha}_t|\boldsymbol{\alpha}_{t+1}, \hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t_2})$$
$$= \frac{p(\boldsymbol{\alpha}_{t+1}|\boldsymbol{\alpha}_t)p(\boldsymbol{\alpha}_t|\hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t}, \boldsymbol{\alpha}_{t_1})}{p(\boldsymbol{\alpha}_{t+1}|\hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t}, \boldsymbol{\alpha}_{t_1})}$$
$$= \mathcal{N}(\boldsymbol{\alpha}_t|\tilde{m}_{\boldsymbol{\alpha}}, \tilde{P}_{\boldsymbol{\alpha}}), \qquad (25)$$

where

$$\tilde{m}_{\boldsymbol{\alpha}} = \mu^{\boldsymbol{\alpha}}_{t|t} + P^{\boldsymbol{\alpha}}_{t|t}A^T(AP^{\boldsymbol{\alpha}}_{t|t}A^T + C)^{-1}(\boldsymbol{\alpha}_{t+1} - A\mu^{\boldsymbol{\alpha}}_{t|t}),$$
$$\tilde{P}_{\boldsymbol{\alpha}} = P^{\boldsymbol{\alpha}}_{t|t} - P^{\boldsymbol{\alpha}}_{t|t}A^T(AP^{\boldsymbol{\alpha}}_{t|t}A^T + C)^{-1}AP^{\boldsymbol{\alpha}}_{t|t}.$$

Again, $p(\boldsymbol{\alpha}_t|\hat{\dot{\psi}}_{t_1+1:t}, s_{t_1+1:t}, \boldsymbol{\alpha}_{t_1}) = \mathcal{N}(\boldsymbol{\alpha}_t|\mu^{\boldsymbol{\alpha}}_{t|t}, P^{\boldsymbol{\alpha}}_{t|t})$ is given by a Kalman filter. Now we can propose $\{\boldsymbol{\alpha}^{(i)}_{t_1+1:t2}\}_{1 \le i \le N}$ using the same backward sampling procedure.

REFERENCES

[1] R. Harle, "A Survey of Indoor Inertial Positioning Systems for Pedestrians," in IEEE Communications Surveys & Tutorials, vol. 15, no. 3, pp. 1281-1293, Third Quarter 2013.

[2] H. Liu, H. Darabi, P. Banerjee and J. Liu, "Survey of Wireless Indoor Positioning Techniques and Systems," in IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), vol. 37, no. 6, pp. 1067-1080, Nov. 2007.

[3] C. Rohrig and S. Spieker, "Tracking of transport vehicles for warehouse management using a wireless sensor network," 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, Nice, 2008, pp. 3260-3265.

[4] B Ristic, S Arulampalam and N Gordon, Ristic, Branko, Sanjeev Arulampalam, and Neil Gordon. "Beyond the Kalman filter: Particle filters for tracking applications", Artech house, 2003.

[5] F. Zhao, S. S. Ge and W. He, "Multiple model based unscented particle filter algorithm for a SINS/CNS integrated navigation system," 2017 American Control Conference (ACC), Seattle, WA, 2017, pp. 1880-1885.

[6] D. Briese, H. Kunze and G. Rose, "UWB localization using adaptive covariance Kalman Filter based on sensor fusion," 2017 IEEE 17th International Conference on Ubiquitous Wireless Broadband (ICUWB), Salamanca, 2017, pp. 1-7.

[7] F. Evennou and F. Marx, "Advanced integration of WIFI and inertial navigation systems for indoor mobile positioning," EURASIP J. Appl. Signal Process, January 2006.

[8] X. Rong Li and V. P. Jilkov, "Survey of maneuvering target tracking. Part I. Dynamic models," in IEEE Transactions on Aerospace and Electronic Systems, vol. 39, no. 4, pp. 1333-1364, Oct. 2003.

[9] R. Hostettler and S. Särkkä, "IMU and magnetometer modeling for smartphone-based PDR," 2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN), Alcala de Henares, 2016, pp. 1-8.

[10] S. Godsill and J. Vermaak, "Variable rate particle filters for tracking applications," IEEE/SP 13th Workshop on Statistical Signal Processing, 2005, Bordeaux, France, 2005, pp. 1280-1285.

[11] P. Bunch and S. Godsill, "Dynamical models for tracking with the variable rate particle filter," 2012 15th International Conference on Information Fusion, Singapore, 2012, pp. 1769-1775.

[12] R. A. Best and J. P. Norton, "A new model and efficient tracker for a target with curvilinear motion," in IEEE Transactions on Aerospace and Electronic Systems, vol. 33, no. 3, pp. 1030-1037, July 1997.

[13] M. Pilt, S. Bonnabel and F. Barbaresco, "Tracking the Frenet-Serret frame associated to a highly maneuvering target in 3D," 2017 IEEE 56th Annual Conference on Decision and Control (CDC), Melbourne, VIC, 2017, pp. 1969-1974.

[14] S. Godsill, J. Vermaak, W. Ng and J. F. Li, "Models and Algorithms for Tracking of Maneuvering Objects Using Variable Rate Particle Filters," in Proceedings of the IEEE, vol. 95, no. 5, pp. 925-952, May 2007.

[15] P. Carey, S. Bennett, J. Lasenby, and T. Purnell, "Aerodynamic analysis via foreground segmentation," in Electronic Imaging, vol. 2017, no. 16, pp. 10-14, Jan. 2017.

[16] O. Cappe, S. Godsill and E. Moulines, "An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo," in Proceedings of the IEEE, vol. 95, no. 5, pp. 899-924, May 2007.

[17] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," Statistics and Computing, vol. 10, no. 3, pp. 197-208, 2000.

[18] C. K. Carter, and R. Kohn. "On Gibbs Sampling for State Space Models, Biometrika, vol. 81, no. 3, pp. 541-553, Aug. 1994.

[19] S. Frühwirth-Schnatter, "Data augmentation and dynamic linear models," Journal of time series analysis, vol. 15, pp. 183-202, Mar. 1994.