Calibrating Away Inaccuracies in Ultra Wideband Range Measurements: A Maximum Likelihood Approach

ANTON LEDERGERBER, (Member, IEEE), AND **RAFFAELLO D'ANDREA**, (Fellow, IEEE)

Institute for Dynamic Systems and Control, ETH Zürich, 8092 Zürich, Switzerland

Corresponding author: Anton Ledergerber (antonl@ethz.ch)

ABSTRACT A calibration framework for an ultra-wideband localization system employing inertial measurement units is presented. No external motion capture system or other sensors are required for the calibration procedure. Given a covariance function for the error in the range measurements, a range measurement model based on a Gaussian process is obtained by maximizing the joint likelihood of angular rate, acceleration, and range measurements. The framework is experimentally evaluated, and it is shown how the resulting measurement model, integrated in a standard Kalman filter, can be used for real-time localization on a platform with limited computational resources. The calibration significantly improves the localization accuracy for randomly generated trajectories and different localization system setups.

INDEX TERMS Gaussian processes, maximum likelihood estimation, sensor fusion, ultra wideband technology.

I. INTRODUCTION

Ultra Wideband (UWB) technology enables low-cost and low-power communication and localization with centimeter accuracy using the same radio transceiver chip [1]. These properties make UWB a potential key technology for the Internet of Things. However, integrating or attaching a UWB transceiver to a device poses many challenges for the technology to achieve the aforementioned accuracy. Often UWB antennas with a small form factor, ideally integrated in a printed circuit board, are required for portable devices. These specifications make omnidirectional antenna design a challenging task for wide bandwidth. Furthermore, these antennas are often mounted on structures or devices that reflect the UWB pulses back towards the antenna, causing the reflected signal to overlap and interfere with the line-of-sight pulse. Although UWB systems feature certain inherent robustness to multipath effects thanks to their short pulse duration, they are not immune to them when pulses overlap [2].

In our previous work [3] we showed that these effects can be partially learned and compensated for using a range measurement model based on Gaussian processes. This paper extends our previous work providing more detailed analytical and experimental results, showing

- how to calibrate such a range measurement model without ground truth data,

- how to calibrate UWB module-dependent delays,
- how parameter choices impact the performance of the framework.

The paper is structured as follows: Section II reviews related work. A Gaussian process model for the UWB range measurement error is introduced in Section III. A maximum likelihood approach to train this model with inertial measurement unit (IMU) and UWB data only is presented in Section IV. Implementation details of the framework are given in Section V, where it is also explained how the resulting model can be integrated in a Kalman filter framework. An experimental evaluation of the proposed framework is presented in Section VI and remarks on future work are made in Section VII. Note that while the symbols and notation are introduced with their first appearance, there is also a overview in the Appendix A.

II. RELATED WORK

The error in the range measurements acquired by a UWB localization system depends on many factors. Errors in clock synchronization, errors in estimating the arrival time of an incoming signal, uncompensated internal propagation delays, multipath propagation and non-line-of sight conditions (NLOS) can all affect the accuracy of a range estimate. Although significant amount of research has investigated the
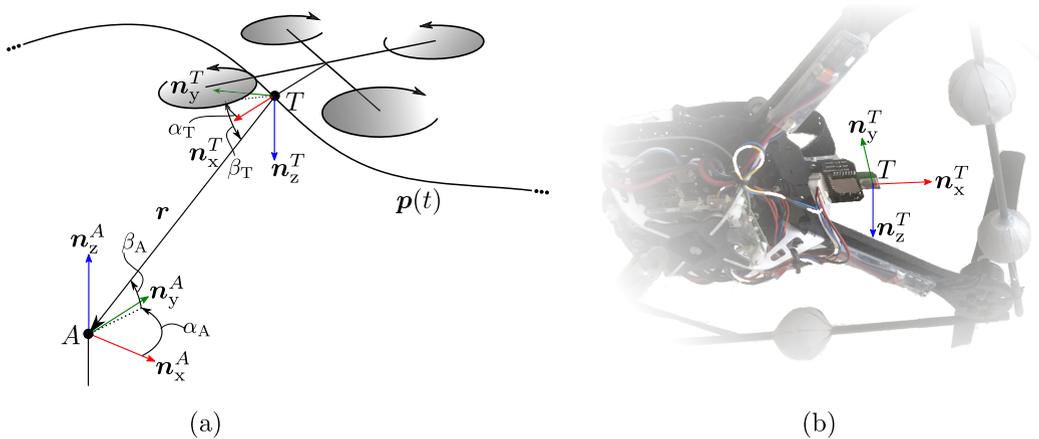
FIGURE 1. A schematic drawing of the setup used for data collection is shown in (a). The coordinate systems of the tag and the anchor are denoted by *T* and *A*, respectively; and the azimuth and elevation angles are denoted by α_A, α_T and β_A, β_T, respectively. The quadcopter platform is pictured in (b). It is visible that UWB pulses from certain directions need to pass the quadcopter body before arriving at the UWB antenna.

effects of obstacles in the far field of the transmitting and receiving antenna (summarized in, for example, [4]), less research has been directed at mitigating the effects of the immediate antenna surrounding, which can cause multipath overlays in the UWB pulse from the direct path in line-of-sight conditions [5]. Additionally, since objects in the near-field of antennas change the antennas' radiation patterns [6], it is generally recommended to keep the near-field free of obstacles. However, this is often not possible in small form-factor, mobile or portable devices [2].

Previous work addressing these real-world problems include [7]–[10]. The effects of angle dependent waveform distortion due to non-ideal antenna design and UWB pulse overlaps on timing estimation are studied in [7], and a new leading edge detection algorithm is proposed that is less sensitive to these effects. If the same antenna in the same immediate surrounding is used, these variations can be compensated for by building a model for the resulting range measurement error. This is done in [8] where a lookup table is used to correct for angle and signal strength-dependent ranging errors. More recently, machine learning techniques have been used to model the correlation between the received UWB pulse waveform and the resulting error in the range measurement. Both [9] and [10] use properties of the received waveform, such as the pulse rise time or the signal strength as features for machine learning algorithms modeling the range measurement error.

However, to build and verify such models for the range measurement error, labeled training data is often required which may not always be available. The work presented here shows how such a model can be obtained without ground truth data. Instead it adapts a calibration procedure shown to work well for combined camera and IMU parameter estimation [11] and extends it to a more complex UWB range model with correlated measurement errors and varying noise variance (heteroscedastic noise [12, p. 273]).

III. RANGE MEASUREMENT MODEL

The most often used, simple range measurement model assumes that the measured range z_r between a mobile UWB module (hereafter known as the tag) and a fixed UWB module (hereafter known as an anchor) is the exact distance, corrupted by zero mean, normally distributed measurement noise η_r with variance σ_r^2. In the following, let the range vector pointing from the tag antenna to the anchor antenna be denoted by $\boldsymbol{r}$ (see Fig. I). Furthermore, let $\mathcal{Z}_r$ denote the set of all range measurements and let the values of the measured range, the range vector and the noise for a specific range measurement $i \in \mathcal{Z}_r$ be denoted by $z_{r,i}$, $\boldsymbol{r}_i$ and $\eta_{r,i}$, respectively. Hence, the standard range measurement model is given as

$$z_{r,i} = \|\boldsymbol{r}_i\| + \eta_{r,i}, \tag{1}$$

where $\|\cdot\|$ denotes the Euclidean norm.

A. RANGE MEASUREMENT ANALYSIS

To assess the suitability of such a measurement model for a real-world application, in which a small UWB-antenna is integrated in a mobile device or robot, the following measurements were taken. A UWB module integrated on a quadcopter platform serving as the tag constantly ranged to an anchor module at a fixed position on the ground. By commanding the quadcopter to yaw around its thrust axis while hovering at different positions, range measurements over different relative poses of the two antennas to each other were collected along with ground truth data provided by a motion capture system. This procedure was repeated three times employing different UWB modules. More details on this setup, depicted in Fig. I, are provided in Section V.

The range measurement error e, i.e. the difference between the exact distance and the measured range

$$e = \|\boldsymbol{r}\| - z_r, \tag{2}$$
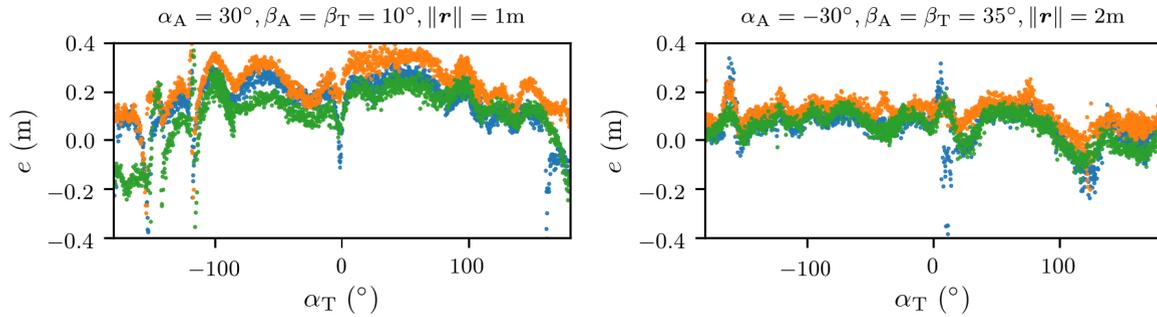
FIGURE 2. The error in range measurements taken over different relative poses is shown (see Fig. 1 for a definition of the angles). Measurements using different UWB module pairs are shown in different colors.

is shown in Fig. 2 for selected relative poses. Since the ranging antennas' z-axes were aligned during the experiments, the relative poses are fully described by the range $\|r\|$, the azimuth α_T and elevation β_T angles of the tag, and the azimuth α_A and elevation β_A angles of the anchor. It is evident that the error in the range measurement is systematic and should not be modelled as white noise. The figure further reveals that the error is correlated with the relative pose of the two ranging modules. This correlation is attributed to the non-omnidirectional radiation pattern of the UWB antenna and to the fact that reflected UWB pulses from the antenna's immediate surrounding might overlap with the UWB pulse from the direct signal path. Given a UWB pulse width of roughly 2 ns [13], such overlaps exist if the indirect path is less than 0.6 m longer than the direct path. With the antenna fixed on a mobile device or robot, these reflections are a functions of the direction of the incoming UWB pulse.

B. RANGE MEASUREMENT MODELING WITH A GAUSSIAN PROCESS

The previously described correlation between the relative pose of the two ranging antennas and the error in the range measurement can be captured by a Gaussian process which assumes that the measured ranges are noisy observations of a unknown function $f(r)$ having the range vector r expressed in the tag frame T (see Fig. 1) as input argument, i.e.

$$z_{r,i} = f(r_i) + \eta_{r,i}. \qquad (3)$$

A Gaussian process for this function is completely specified by its mean $m(r)$ and covariance function $k(r, r')$ [14, p. 13], i.e.

$$m(r) = \mathbb{E}[f(r)] \qquad (4)$$
$$k(r, r') = \mathbb{E}\left[(f(r) - m(r))\left(f(r') - m(r')\right)\right]. \qquad (5)$$

The mean function is chosen to be the expected value of the standard range measurement model

$$m(r) = \|r\|. \qquad (6)$$

We later discuss in Section III-C how to adapt the mean function in case of constant, UWB module specific offsets.

The covariance function is chosen to be

$$k(r, r') = \theta_0 \exp\left(-\underbrace{\frac{1 - \frac{r^\mathsf{T} r'}{\|r\|\|r'\|}}{\theta_1}}_{\text{angle term}} - \underbrace{\left(\frac{\|r\| - \|r'\|}{\theta_2}\right)^2}_{\text{distance term}}\right), \qquad (7)$$

where $\theta_0, \theta_1, \theta_2$ are adjustable parameters. Given two range measurements $i, j \in \mathcal{Z}_r$, this covariance function models the similarity of ranging errors e_i and e_j as a function of the angle between the incoming UWB pulses to the tag (angle term) and the difference in the distance the UWB pulses have to travel (distance term), i.e.

$$e_i = \|r_i\| - z_{r,i} \qquad (8)$$
$$\text{cov}(e_i, e_j) = k(r_i, r_j) + \delta_{ij}\sigma_r^2, \qquad (9)$$

where δ_{ij} is the Kronecker delta which is one iff $i = j$ and zero otherwise. Hence, if range measurements i, j were made in roughly the same direction, the angle term of (7) would be small. If they were also made over roughly the same distances, the distance term of (7) would also be small, resulting in $k(r_i, r_j) \approx \theta_0$. To the contrary, if either the directions or the distances of the two measurements were very different, the covariance function value would only be a fraction of θ_0. This behavior is defined by the values of θ_1 and θ_2, and is visualized in Fig. 3. We show in Appendix B that $k(r, r')$ is a valid covariance function.

Note that the range vector r is not a full parameterization of the relative pose. Nonetheless the previously defined Gaussian process for the function $f(r)$ is able to describe the measured range well, as shown in Section VI. Other Gaussian processes for functions having the complete relative pose as input arguments, or for functions describing the ranging error with properties of the received waveform as used in [10], are not investigated in this work. Nonetheless, the framework presented hereafter is applicable to any Gaussian process describing the range measurement model.

The application of a standard Gaussian process can be computationally demanding when using many training points. To lower the computational demand, the sparse Gaussian process using pseudo input points (SPGP) is employed. This sparse approximation to a Gaussian Process was first presented in [15] and is now termed
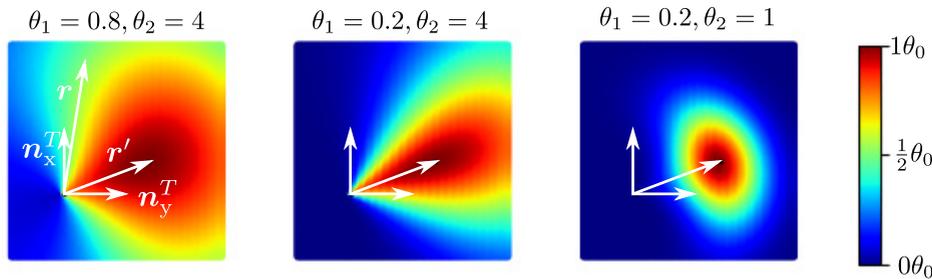
FIGURE 3. The value of the covariance function $k(r, r')$ given in (7) is shown as a function of $r \in \text{span}(n_x^T, n_y^T)$ while keeping r' fixed. For example, for r as explicitly shown in the first plot, it evaluates to $k(r, r') = \frac{1}{2}\theta_0$. For purpose of comparison, the three plots were made with different values of the kernel parameters θ_1 and θ_2 — in the framework described in Section IV, these parameters are optimized to maximize the data likelihood.

fully-independent training conditional sparse Gaussian process [16]. A complete introduction to the SPGP is beyond the scope of this paper; therefore, only the necessary preliminaries for further discussion are provided and the reader is referred to [15] and [16] for more details. While a standard Gaussian process model directly uses its N training data to infer the posterior distribution over the measured range $z_{r, \text{test}}$ for the test range vector r_{test}, the SPGP performs the inference via M latent variables (where $M \ll N$), termed pseudo input points. In the following we denote the mean and the variance of this posterior distribution given by the SPGP with $\text{SPGP}_m(r_{\text{test}})$ and $\text{SPGP}_{\sigma^2}(r_{\text{test}})$, respectively. This distribution is fully defined by the SPGP's mean and covariance function with parameters θ_0, θ_1 and θ_2, its training data, which consist of N range vectors $\Gamma = [r_0, r_1, \ldots, r_{N-1}]$ and the corresponding measured ranges $z_r = [z_{r,0}, z_{r,1}, \ldots, z_{r,N-1}]$, and the location of its M pseudo input points $\bar{\Gamma} = [\bar{r}_0, \bar{r}_1, \ldots, \bar{r}_{M-1}]$. A nice property of the SPGP is that input-dependent noise can be modeled by altering the location of the pseudo input points.

When comparing the error characteristics in Fig. 2 for line-of-sight and non-line-of-sight conditions (non-line-of-sight: $\beta_T = 10$ deg and $\alpha_T \notin [-140 \text{ deg}, 140 \text{ deg}]$), it is visible that the error is less repeatable for range measurements in non-line-of sight conditions. The lack of confidence in predicting the measured range for such regions can be modeled by input-dependent noise when employing the SPGP as previously mentioned.

The parameters of the SPGP $\Theta := (\bar{\Gamma}, \theta_0, \theta_1, \theta_2)$ are found by maximizing the likelihood of the training data. In the following, we denote with $K_{\Gamma\Gamma}$, $K_{\bar{\Gamma}\bar{\Gamma}}$, and $K_{\Gamma\bar{\Gamma}}$ the covariance matrices, whose entries are given by evaluating the covariance function for all pairs of training points, pseudo input points, and combinations of training and pseudo input points, respectively. Denoting with m the vector of stacked mean function evaluations for the training range vectors, the likelihood of the range measurements is given as [15]

$$p_\Theta(z_r | \Gamma) = \mathcal{N}(z_r | m, K_{\Gamma\bar{\Gamma}} K_{\bar{\Gamma}\bar{\Gamma}}^{-1} K_{\Gamma\bar{\Gamma}}^T + \Lambda + \sigma_r^2 I), \quad (10)$$

where $\Lambda = \text{diag}(K_{\Gamma\Gamma} - K_{\Gamma\bar{\Gamma}} K_{\bar{\Gamma}\bar{\Gamma}}^{-1} K_{\Gamma\bar{\Gamma}}^T)$ and where I is the identity matrix of size N. Maximizing this likelihood with

respect to the parameters Θ moves the location of the pseudo input points $\bar{\Gamma}$ outside regions with high variance, hence to a certain extent it can model input dependent noise.

C. EXTENSION FOR MODULE SPECIFIC OFFSETS
Small manufacturing and assembly differences of the UWB modules can cause differences in internal propagation delays, resulting in offsets in the range measurements. These offsets can vary from module to module and are generally identified in a separate calibration routine as outlined in [17]. For example, in Fig. 2 the range measurement error shown in orange appears to have a constant offset to the error of the other two datasets obtained with different UWB module pairs. To accommodate these module specific offsets when working with multiple anchors, the mean function of the Gaussian Process given in (6) can be replaced by

$$m(r) = \|r\| - o, \quad (11)$$

where o denotes the constant offset associated to the tag-anchor module pair with which a range measurement is taken. The offsets for all module pairs o are included in the parameter set of the SPGP, $\Theta := (\bar{\Gamma}, \theta_0, \theta_1, \theta_2, o)$.

IV. MODEL LEARNING
To train the range measurement model based on a Gaussian process presented in Section III-B, labeled training data (Γ, z_r) is necessary. Historically, this required UWB localization system developers to collect ground truth data using a more accurate system, such as a infrared motion capture system. However, UWB localization technology aims to replace other more expensive systems. Using these other systems in a calibration routine is therefore undesirable. In contrast, consumer grade inertial measurement units (IMUs) are cheap and can be found in most mobile devices. Furthermore, they are often integrated with tag devices.

This section shows how calibration of a UWB range measurement model can be performed using only UWB range and inertial measurements adapting the continuous-time batch optimization framework presented in [18] and used in [11] to calibrate a camera-IMU setup.

A. THE MAXIMUM A POSTERIORI ESTIMATE

Consider a tag equipped with an IMU and a UWB module moving through a space equipped with anchors at known locations while collecting accelerometer, angular rate and UWB range measurements. The sets of these measurements are denoted by $\mathcal{Z}_{\text{acc}}$, $\mathcal{Z}_{\text{gyr}}$ and $\mathcal{Z}_{\text{r}}$, respectively. Let $\boldsymbol{x}(t) = (\boldsymbol{p}(t), R(t), \boldsymbol{b}_{\text{acc}}(t), \boldsymbol{b}_{\text{gyr}}(t))$ denote the tag's state over the entire experiment, described by the position trajectory of the tag antenna $\boldsymbol{p}(t)$ given in the inertial frame, the rotation $R(t)$ from the inertial to tag frame T, and the trajectories of the accelerometer $\boldsymbol{b}_{\text{acc}}(t)$ and gyroscope $\boldsymbol{b}_{\text{gyr}}(t)$ biases.

Using Bayes' rule, the maximum a posteriori estimate of the trajectory is written as

$$
\begin{aligned}
&p_\Theta\left(\boldsymbol{x}(t)|\mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}, \mathcal{Z}_{\text{r}}\right) \\
&= \frac{p_\Theta(\mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}, \mathcal{Z}_{\text{r}}|\boldsymbol{x}(t))p_\Theta(\boldsymbol{x}(t))}{p_\Theta(\mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}, \mathcal{Z}_{\text{r}})}.
\end{aligned} \tag{12}
$$

Assuming that the sensor measurements are conditionally independent given $\boldsymbol{x}(t)$, and assuming that only the range measurements are dependent on the parameters Θ, the posterior is rewritten as

$$
\begin{aligned}
&p_\Theta\left(\boldsymbol{x}(t)|\mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}, \mathcal{Z}_{\text{r}}\right) \\
&= \frac{p(\mathcal{Z}_{\text{gyr}}|\boldsymbol{x}(t))p(\mathcal{Z}_{\text{acc}}|\boldsymbol{x}(t))p_\Theta(\mathcal{Z}_{\text{r}}|\boldsymbol{x}(t))p(\boldsymbol{x}(t))}{p(\mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}, \mathcal{Z}_{\text{r}})}.
\end{aligned} \tag{13}
$$

The maximum a posteriori estimate of the state trajectory $\boldsymbol{x}(t)$ and the parameters Θ can be found by minimizing the negative logarithm of the posterior

$$
\begin{aligned}
&\{\boldsymbol{x}^*(t), \Theta^*\} \\
&= \underset{\boldsymbol{x}(t),\Theta}{\arg\min} \, (-\log\left(p(\mathcal{Z}_{\text{acc}}|\boldsymbol{x}(t))\right) - \log\left(p(\mathcal{Z}_{\text{gyr}}|\boldsymbol{x}(t))\right) \\
&\quad - \log\left(p(\mathcal{Z}_{\text{r}}|\boldsymbol{x}(t), \Theta)\right) - \log\left(p(\boldsymbol{x}(t))\right)),
\end{aligned} \tag{14}
$$

where the term $\log\left(p\left(\mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}, \mathcal{Z}_{\text{r}}\right)\right)$ was not considered as it does not influence the optimized values. All terms on the right hand side of (14) are specified in the following two subsections.

B. THE MEASUREMENT LIKELIHOODS

Assuming that the accelerometer measurements are conditionally independent given $\boldsymbol{x}(t)$ and have a Gaussian distribution, their likelihood is given by

$$
p(\mathcal{Z}_{\text{acc}}|\boldsymbol{x}(t)) = \prod_{k \in \mathcal{Z}_{\text{acc}}} p(z_{\text{acc},k}|\boldsymbol{x}(t_k)) \tag{15}
$$

$$
p(z_{\text{acc},k}|\boldsymbol{x}(t_k)) = \mathcal{N}(z_{\text{acc},k}|h_{\text{acc}}(\boldsymbol{x}(t_k)), \Sigma_{\text{acc}}), \tag{16}
$$

where $z_{\text{acc},k}$ denotes the value of the accelerometer measurement $k \in \mathcal{Z}_{\text{acc}}$ and where $\boldsymbol{x}(t_k)$ is used to denote the tag's state at the time measurement k was taken. The covariance matrix of the accelerometer measurement noise is denoted by Σ_{acc}, and h_{acc} is used to denote the deterministic part of the accelerometer measurement model which is further specified in Appendix C. The likelihood of the angular rate measurements can be calculated similarly and their measurement model is also given in Appendix C.

The range measurements $\mathcal{Z}_{\text{r}}$ are not conditionally independent given $\boldsymbol{x}(t)$, but are correlated according to the covariance function given in Section III-B. In order to find their likelihood given $\boldsymbol{x}(t)$, the range vectors $\boldsymbol{r}_i$, $i \in \mathcal{Z}_{\text{r}}$ need to be expressed as function of the state $\boldsymbol{x}(t)$, i.e.

$$
\boldsymbol{r}_i(\boldsymbol{x}(t_i)) = R^{\mathsf{T}}(t_i)\,(\boldsymbol{a}_i - \boldsymbol{p}(t_i)) \quad \forall i \in \mathcal{Z}_{\text{r}}, \tag{17}
$$

where $\boldsymbol{p}(t_i)$ and $R(t_i)$ denote the position and the attitude of the tag when the range measurement i is obtained, and where $\boldsymbol{a}_i$ denotes the position of the anchor module with which the range measurement i is taken. According to (10), the likelihood of the range measurements is then given as

$$
p_\Theta(\mathcal{Z}_{\text{r}}|\boldsymbol{x}(t)) = \mathcal{N}\left(\boldsymbol{z}_{\text{r}}|\boldsymbol{m}, K_{\Gamma\bar{\Gamma}}K_{\bar{\Gamma}\bar{\Gamma}}^{-1}K_{\Gamma\bar{\Gamma}}^{\mathsf{T}} + \Lambda + \sigma_{\text{r}}^2 I\right) \tag{18}
$$

where $\boldsymbol{m}$, $K_{\Gamma,\bar{\Gamma}}$ and Λ are all calculated as a function of the state $\boldsymbol{x}(t)$ via (17), and where $\boldsymbol{z}_{\text{r}}$ is the vector of stacked measured ranges $z_{\text{r},i}$, $i \in \mathcal{Z}_{\text{r}}$.

C. THE PRIOR

As outlined in [18], if a model for the tag's state evolution is at hand, it can be used to specify a prior on the state trajectory. Given the evolution function $\Phi(\boldsymbol{x}(t))$ and continuous-time, additive, Gaussian process noise with covariance Q, the tag's state evolves as follows:

$$
\dot{\boldsymbol{x}}(t) = \Phi(\boldsymbol{x}(t)) + \boldsymbol{w}(t). \tag{19}
$$

Using this model, the prior on the state trajectory is given as

$$
p(\boldsymbol{x}(t)) \propto p(\boldsymbol{x}(t_{\text{start}})) \exp\left(-\frac{1}{2}\int_{t_{\text{start}}}^{t_{\text{end}}} \boldsymbol{d}^{\mathsf{T}}(\tau)Q^{-1}\boldsymbol{d}(\tau)d\tau\right), \tag{20}
$$

where $\boldsymbol{d}(t) = \dot{\boldsymbol{x}}(t) - \Phi(\boldsymbol{x}(t))$ and where $p(\boldsymbol{x}(t_{\text{start}}))$ denotes the prior belief of the initial state which is assumed to be Gaussian-distributed. If the ratio of the number of optimization variables to the number of measurements is large, such a prior becomes important as it can serve as a regularizer [18].

D. THE POSTERIOR RANGE MEASUREMENT MODEL

Given the optimized parameters Θ^*, the measured ranges $z_{\text{r},i}$, $i \in \mathcal{Z}_{\text{r}}$ and the range vectors calculated with the maximum a posteriori state estimate $\boldsymbol{r}_i(\boldsymbol{x}^*(t_i))$, $i \in \mathcal{Z}_{\text{r}}$, the SPGP defines a posterior distribution over the function $f(\boldsymbol{r})$ as mentioned in Section III-B. Hence the range measurement model using this SPGP is given as

$$
z_{\text{r},i} = \text{SPGP}_m(\boldsymbol{r}_i) + \sqrt{\text{SPGP}_{\sigma^2}(\boldsymbol{r}_i)}\eta_1, \tag{21}
$$

where η_1 is zero mean additive Gaussian noise with unit variance. The computational cost to evaluate this range measurement model is $\mathcal{O}(M)$ for the mean and $\mathcal{O}(M^2)$ for the variance using an SPGP with M pseudo input points [15].

V. IMPLEMENTATION AND INTEGRATION IN A KALMAN FILTER FRAMEWORK

The previously described framework to calibrate an UWB range measurement model was implemented and experimentally evaluated. This section provides implementation details and shows how the extracted range measurement model can be integrated in a Kalman filter framework, which can be run in real-time on a device with limited computational power.

A. B-SPLINE REPRESENTATION OF THE CONTINUOUS TIME STATE TRAJECTORY

The continuous-time state trajectory $x(t)$ is expressed using standard, uniform cubic B-splines for the position trajectory $p(t)$ expressed in the inertial frame and using cumulative, uniform cubic B-splines [19], [20] for the quaternion $q(t)$ describing the rotation $R(t)$ from the inertial frame to the tag frame T. The biases of the accelerometer $b_{acc}(t)$ and the gyroscope $b_{gyr}(t)$ are both modelled using standard, uniform cubic B-splines.

The spacing of the knots for the position and quaternion splines is set to 50 ms and the splines describing the biases of the IMU have a spacing of 5 s.

B. SOFTWARE AND HARDWARE

The optimization routine is implemented in Tensorflow [21]. This permits easy use of the computational power of the GPU and Tenserflow's automatic differentiation facilitated the implementation.

The UWB modules used on the tag and the anchors were DWM 1000 modules [22] configured with the settings given in Table 1. The anchor microcontroller is a STM32f4 and the tag microcontroller is a Snapdragon flight board. The tag module ranges with the anchors at 200 Hz by employing the two-way ranging algorithm with repeated reply described in [23].

TABLE 1. The settings of the DWM1000 modules used for ranging.

Channel	5
Pulse Repetition Frequency	16 MHz
Data Rate	6.8 Mbps
Preamble Length	128 Symbols
Output Power	41.3 dBm/MHz

The MPU9250 IMU module, integrated on the snapdragon flight board, provides angular rate and acceleration measurements with a frequency of 1 kHz.

The quadcopter platform serving as the tag uses the frame, motors and motor controllers of an Ascending Technology Hummingbird and is equipped with the DWM 1000 module and the Snapdragon flight board. The frame is made of carbon fiber and some magnesium rods which both reflect UWB pulses to a certain extent.

C. INTEGRATION IN A KALMAN FILTER FRAMEWORK

The optimization procedure given in (14) allows us to estimate the state of a rigid body given a batch of IMU and range measurements. On platforms with limited computational power such as wearables, this optimization cannot be performed in real-time. However, using one of the methods outlined in [24], the range measurement model obtained through the batch optimization given in (21) can be integrated in a Kalman filter which recursively estimates the state.

To investigate the gain in localization accuracy when employing the calibrated range measurement model, it is used to update a Kalman filter which estimates the tag's position p, its velocity v, both expressed in the inertial frame, and its orientation with respect to the inertial frame $R \in SO(3)$. By using the acceleration z_{acc} and the angular rate z_{gyr} measured with the IMU as inputs to the filter, its state evolves as

$$\dot{p} = v \tag{22}$$
$$\dot{v} = Rz_{acc} + g \tag{23}$$
$$\dot{R} = R[\![z_{gyr}]\!]_\times, \tag{24}$$

where g is the gravitational acceleration expressed in the inertial frame and where $[\![z_{gyr}]\!]_\times$ denotes the matrix form of the cross product, defined such that $[\![z_{gyr}]\!]_\times b = z_{gyr} \times b$ for all $b \in \mathbb{R}^3$. The measurement noise of the IMU is encoded as the process noise using the standard Kalman filter formulation. While the prediction step of this filter is performed using an extended Kalman filter update step, the measurement update step is performed using the unscented transform [24].

VI. EXPERIMENTAL RESULTS

All experiments were performed in the Flying Machine Arena of ETH Zurich [25]. This arena is equipped with a motion capture system providing ground truth data for all experiments with an accuracy of the order of millimeters and milliradians at 200 Hz.

A. DATA COLLECTION

A quadcopter equipped with an IMU and a UWB module autonomously flew random trajectories while constantly logging data. It was fusing the measurements in the state estimator presented in Section V-C, but used the standard UWB range measurement model given in (1). Its state estimate served as an initial guess for the optimization routine
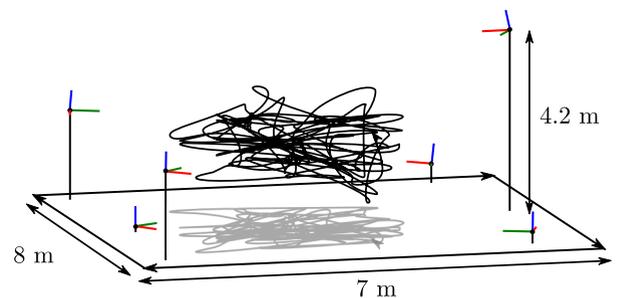
FIGURE 4. The anchors' coordinate systems and the random trajectory flown during data collection are shown. Note that the quadcopter was yawing around its thrust axis while taking range measurements to all depicted anchors.
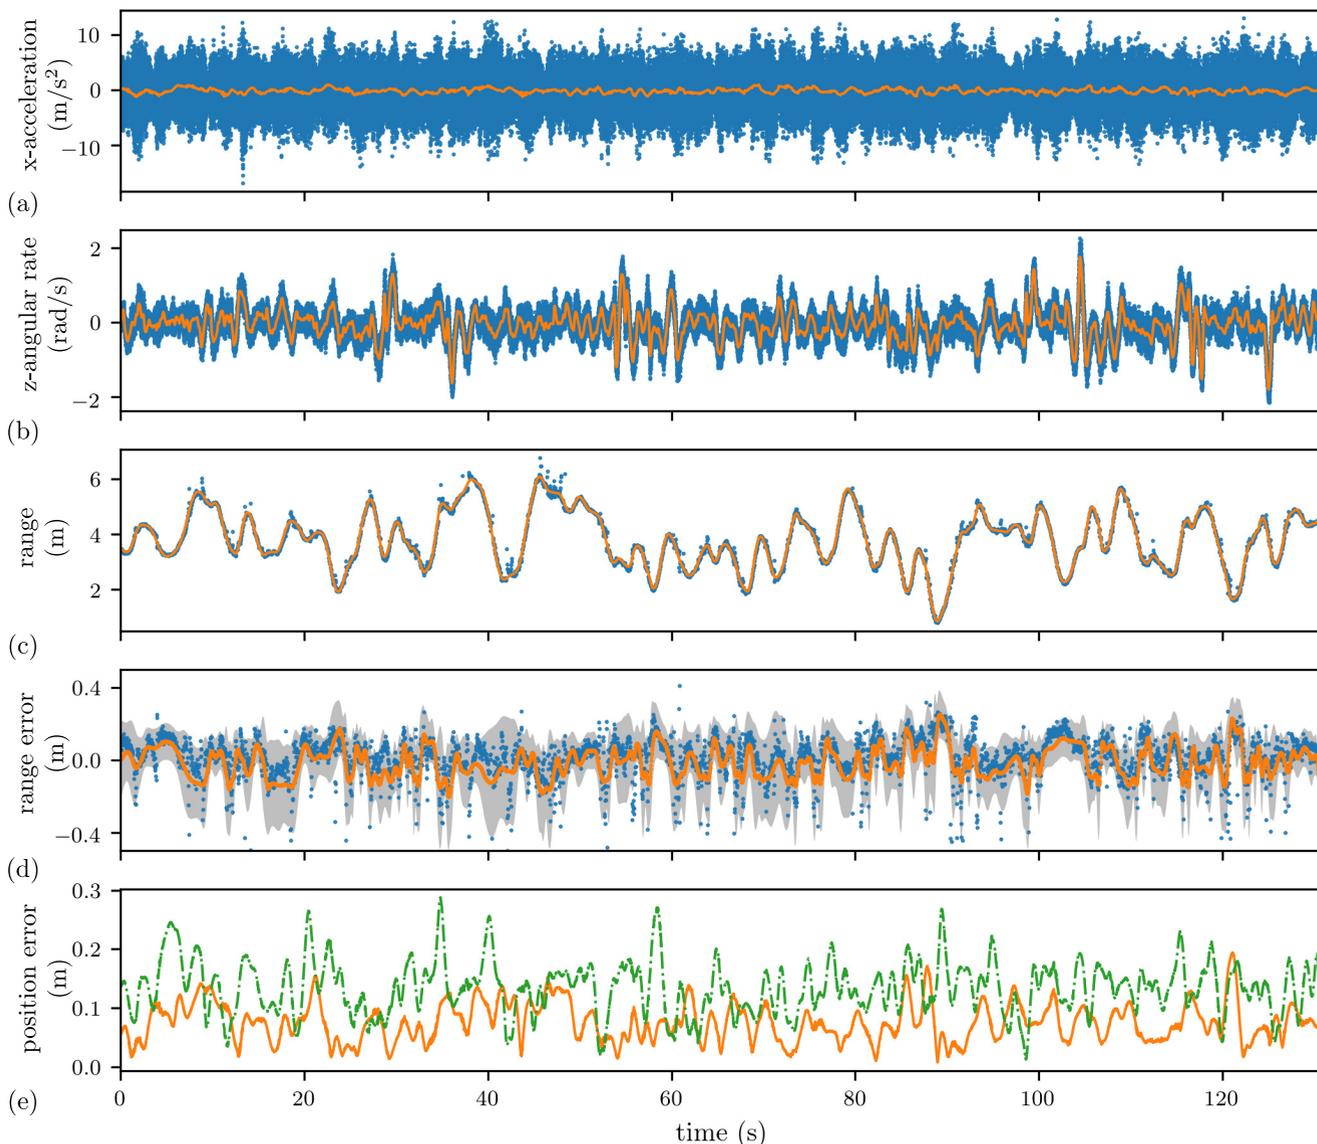
FIGURE 5. A selection of the measured (blue) and the fitted (orange) quantities for the trajectory visualized in Fig. 4 are shown in plots (a) to (d); (a) the x-axis value of z_{acc} and $h_{acc}(x*(t))$, (b) the z-axis value of z_{gyr} and $h_{gyr}(x*(t))$, (c) z_r and $SPGP_m(r*(x*(t)))$ obtained with one anchor, (d) e and $e_{SPGP}(r(x*(t)))$ along with the predicted uncertainty bounds $\pm 2\sqrt{SPGP_{\sigma^2}(r(x*(t)))}$ for one anchor. Plot (e) shows the position error $\|p(t) - p*(t)\|$ of the trajectory optimized with the proposed range measurement model (solid orange) and of the trajectory optimized with the standard range measurement model (green dash-dotted).

given in (14) and was logged in addition to the IMU and range measurements. The anchor placement and the random trajectory flown are depicted in Fig. 4. The flight time was approximately 130 s resulting in about 130 000 IMU measurements and 25 000 range measurements.

B. MODEL FIT
The optimization of (14) was performed with the collected data. The tag's state trajectory $x(t)$ was initialized with the state estimated during data collection. No prior as outlined in Section IV-C was given on this state trajectory, except for the bias states for which a first order random walk model was assumed.

The number of pseudo input points was $M = 30$ and all initial values for the parameters $\bar{\Gamma}, \theta_0, \theta_1$ and θ_2 were randomly initialized as further discussed in Section VI-D. The standard deviation of the range measurement was fixed at $\sigma_r = 0.05$ m and all anchor-tag module specific biases o were initialized with zero.

Figure 5 shows the fit of the acceleration along the x-axis in plot (a) and the fit of the angular rate around the z-axis in plot (b). The fit of the other axes look similar. Even though the accelerometer measurements are very noisy due to the vibrations induced by the spinning motors, these IMU measurements prove to be sufficient to learn a range measurement model. The fit of the predicted measured range
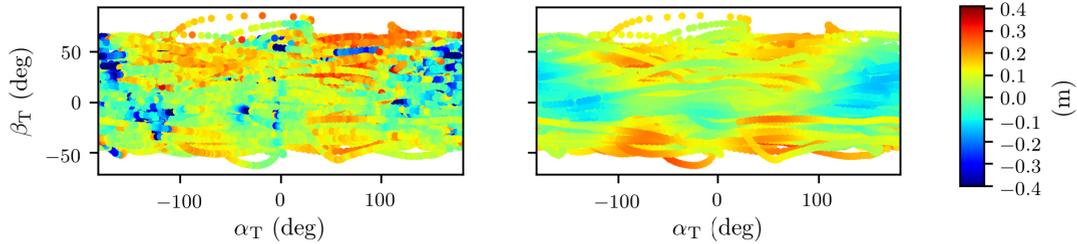
FIGURE 6. The left plot shows the the range measurement error e calculated with the ground truth measurements of the motion capture system for the experiment shown in Fig. 4. The error is plotted over the azimuth α_T and elevation β_T of the tag antenna (see Fig. I for a definition of these angles). The right plot shows the range measurement error predicted by the model $e_{SPGP}(r)$, which was obtained using only inertial and UWB range measurements. Note that the range and the module biases also influence the predicted error aside from the azimuth α_T and elevation β_T angle. This explains the discontinuities visible in the right plot.
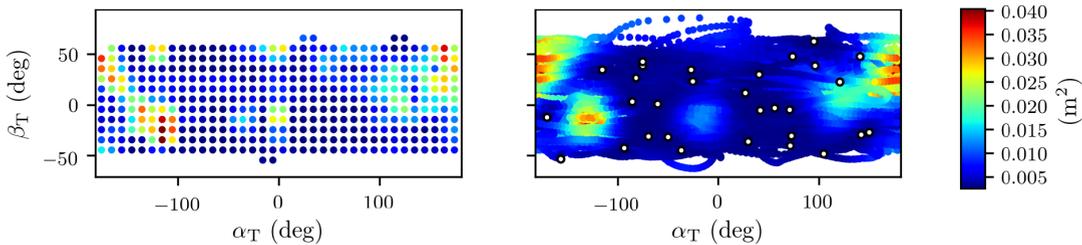
FIGURE 7. The variance of the residual range measurement error e_{res} (left), calculated by binning the measurements based on the azimuth and elevation angle, and the predicted variance (right) plotted over the azimuth α_T and elevation β_T of the tag antenna (see Fig. I for a definition of these angles). The white dots are the locations of the pseudo input points $\bar{\Gamma}$ projected into the azimuth-elevation plane. The NLOS regions clearly show a higher variance.

given by this model, i.e. $SPGP_m(r)$, is shown in plot (c) for one anchor. Plot (d) shows the fit of the associated predicted error in the measured range e_{SPGP} given as

$$e_{SPGP} = \|r\| - SPGP_m(r) \qquad (25)$$

to the errors e calculated with the ground truth data. The uncertainty bounds $\pm 2\sqrt{SPGP_{\sigma^2}(r)}$ are shown in gray. The standard deviation of the error in the range measurements e to all anchors is 0.13 m while the standard deviation of the residual error e_{SPGP}

$$e_{res} = e - e_{SPGP} \qquad (26)$$

is 0.11 m. The last plot (e) in Fig. 5 shows the error in the optimized position trajectory $\|p(t) - p^*(t)\|$. The root mean squared error (RMSE) of this trajectory is 0.08 m. For comparison, the error in the position of a trajectory optimized with the standard range measurement model as given in (1) is shown in green. The RMSE of this second trajectory is 0.14 m. Hence, the rather modest reduction of the standard deviation of the range measurement error by a factor of $0.11/0.13 = 0.8$ concurs with a reduction in the RMSE of the position by a factor of $0.08/0.14 = 0.6$. Given Gaussian range error distributions, one would expect these factors to be equal [26]. However, as described in Section III-A, the range error does not show a Gaussian distribution, instead it is correlated and has outliers in certain regions. This might explain the larger improvement in the RMSE of the position with the range measurement model given in (3).

The actual and the predicted error for all range measurements is plotted over the azimuth α_T and elevation angle β_T in Fig. 6. Note that the discontinuities in the predicted error

arise from the different distances over which the range measurements were taken which also influence the range measurement error along with the direction of the incoming UWB signal. It is visible that the optimization procedure given in (14) was able to extract an accurate range measurement error model without using ground truth data.

Looking at the predicted variance and the location of the pseudo input points in Fig. 7, it is visible how the SPGP can model heteroscedastic noise by moving the pseudo input points away from regions with a lot of noise or model uncertainty. The left plot in this figure shows the variance of binned range measurements based on their azimuth and elevation angle. The region with non-line-of-sight conditions clearly shows a higher predicted and measured variance. Also the region along the antenna x-axis, where the antenna gain is smallest [22], shows a higher variance. The model obtained from the optimization procedure also seems to capture these regions.

C. MODEL APPLICATION TO OTHER DATASETS

To see how well the obtained range measurement model generalizes to other trajectories and setups, it was applied to the following four datasets including the dataset used for training to investigate over fitting.

- Dataset 1: The training dataset of the UWB range measurement model
- Dataset 2: A dataset recorded with the same anchor setup, but with a different trajectory
- Dataset 3: A dataset recorded with the same anchor setup, but with a different trajectory and a different DWM1000 tag module

- Dataset 4: A dataset recorded with a different anchor setup, with a different trajectory but with the same DWM1000 tag module

The standard deviation of the range measurement error and the standard deviation of the residual range measurement error of all datasets is given in Table 2. Independent of the trajectory flown, the tag's DWM1000 module or the anchor setup, the standard deviation could be reduced by 2 cm. The model obtained was further integrated in a Kalman filter as outlined in Section V-C and tested on the datasets. For comparison, a Kalman filter employing the standard range measurement model as given in (1) was also run on the datasets. The RMSE of their position estimates, denoted by $\text{RMSE}(\hat{\boldsymbol{p}}_{\text{KF, prop.}})$ and $\text{RMSE}(\hat{\boldsymbol{p}}_{\text{KF, std.}})$, respectively, are also shown in the table. Note that the standard deviation of the range measurement noise for the Kalman filter employing the standard model had to be increased from $\sigma_{\text{r}} = 0.05$ m to $\sigma_{\text{r}} = 0.13$ m to compensate for the unmodelled error characteristics. Without this measure, the estimate would diverge.

TABLE 2. Performance on datasets (1 training, 2 different trajectory, 3 different tag module and trajectory, 4 different anchor setup and trajectory), all values are given in meters.

Dataset	1	2	3	4
stdDev(e)	0.13	0.15	0.13	0.14
stdDev(e_{res})	0.11	0.13	0.11	0.12
RMSE($\hat{\boldsymbol{p}}_{\text{KF, std.}}$)	0.16	0.19	0.16, 0.16*	0.19
RMSE($\hat{\boldsymbol{p}}_{\text{KF, prop.}}$)	0.10	0.14	0.14, 0.11*	0.13

Comparing the RMSE of the position estimate obtained with the Kalman filters with the one obtained from the batch optimization for dataset 1 given in Section VI-B, i.e. 0.16 m vs. 0.14 m for the standard model and 0.10 m vs. 0.08 m for the proposed model, it is visible that the batch optimization approach outperforms the recursive Kalman filter approach.

As Dataset 3 was recorded with a different DWM1000 tag module, showing a different module specific delay, the offsets $\boldsymbol{o}$ for tag-anchor module pairs (as discussed in Section III-C) were off by a constant value, which explains the smaller reduction in the RMSE of the position estimate in Dataset 3. However, by comparing the reported range of

the old and the new tag module to a specific anchor, this constant value can be determined and the offsets can be updated. Using the updated offsets, the RMSE of the KF employing the proposed range measurement model could be further reduced while the the RMSE of the KF employing the standard range measurement model stayed approximately the same. The values obtained with the updated offsets are marked with an asterisk $*$ in Table 2.

D. GAUSSIAN PROCESS PARAMETER INITIALIZATION
The optimization of the parameters for Gaussian processes, i.e. $\theta_0, \theta_1, \theta_2, \bar{\Gamma}$ and $\boldsymbol{o}$ in our case, is usually not a convex problem [27]. A common approach to tackle this issue is to run the optimization with multiple starting points sampled from a prior distribution, and choosing the optimized values with the smallest loss at the end.

To investigate how sensitive the optimization procedure was to the number of pseudo input points M, their initial values and the initial values of parameters θ_0, θ_1 and θ_2, the optimization was run multiple times using different initial values for a given number of pseudo input points M. The pseudo input points were randomly sampled from the unit sphere and then multiplied by a sample drawn from the uniform distribution $U_{\|\bar{\boldsymbol{r}}\|}(1, 7)$. The parameters θ_0, θ_1 and θ_2 were sampled from the uniform distributions $U_{\theta_0}(0.1, 1)$, $U_{\theta_1}(0.1, 3)$ and $U_{\theta_2}(1, 10)$. The anchor-tag module specific offsets were always initialized with zero.

Using the optimized models the standard deviation of the residual range measurement error, stdDev(e_{res}), and the RMSE of the position estimate of a KF employing these models, $\text{RMSE}(\hat{\boldsymbol{p}}_{\text{KF, prop}})$, were calculated. The values for datasets 1 and 4 are shown in Fig. 8. While the initial values of the parameters do not seem to play a major role, the number of pseudo input points should not be too large or too small. It is apparent that about $M = 30$ pseudo input points are needed to capture the error characteristics. More points only lead to an improvement in the training data set, but not to an improvement in a dataset obtained with a different anchor setup and trajectory. When using less than $M = 30$ pseudo input points, there is still an improvement in the standard deviation of the
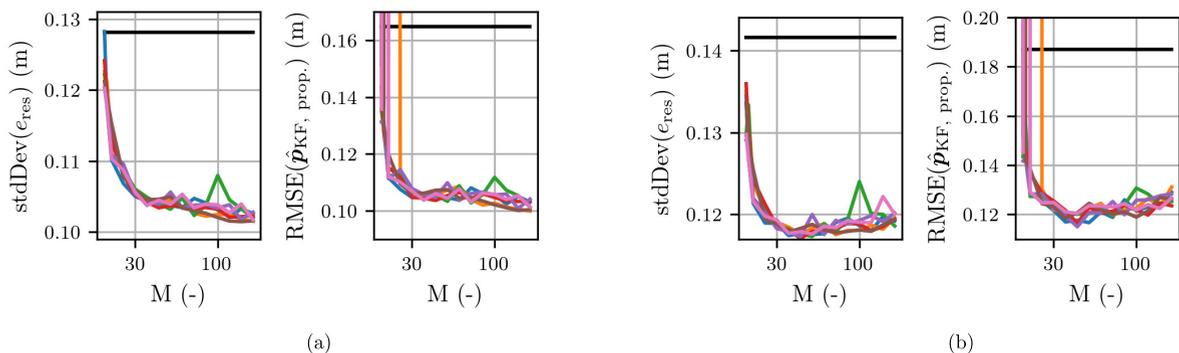
FIGURE 8. The figure shows the effect of random parameter initialization and different number of pseudo input points M on the standard deviation of the residual ranging error and the the RMSE of the position estimate obtained by the Kalman filter employing the model. The black lines show the standard deviation of the range measurement error and the RMSE of the position estimate obtained by a Kalman filter run with the standard range measurement update. (a) Dataset 1. (b) Dataset 4.

residual error, but the position estimate of the Kalman filter sometimes diverges. Looking at these cases, it was visible that the parameters converged to values resulting in a too small predicted variance for regions with a lot of model uncertainty. Hence for the proposed kernel and application, M should be chosen to be between 30 and 70.

TABLE 3. Mathematical symbols in alphabetical order.

$\boldsymbol{a}$	The position of the anchor with which a range measurement is taken
$\boldsymbol{b}_{\text{acc}}(t), \boldsymbol{b}_{\text{gyr}}(t)$	The accelerometer and gyroscope bias over the entire experiment
e	The error in the range measurement
e_{res}	The residual error, i.e. the error in the range measurement not explained by the model
e_{SPGP}	The error of the measured range predicted by the model employing the SPGP
$f(\boldsymbol{r})$	An unknown function describing the range measurement model
$\boldsymbol{g}$	The gravitational acceleration expressed in the inertial frame
h_{acc}	The deterministic part of the accelerometer measurement model.
$k(\boldsymbol{r}, \boldsymbol{r}')$	The covariance function of the Gaussian process
$K_{\Gamma\Gamma}, K_{\Gamma\bar{\Gamma}}, K_{\bar{\Gamma}\bar{\Gamma}}$	The covariance matrices
$m(\boldsymbol{r}), \boldsymbol{m}$	The mean function of the Gaussian process and the vector of stacked mean function evaluations
M, N	The number of pseudo input and training points
$o, \boldsymbol{o}$	A UWB module pair specific offset, and the vector of stacked offsets
$\boldsymbol{p}(t)$	The tag's position over the entire experiment
$\boldsymbol{p}^*(t)$	The fitted position trajectory over the entire experiment
$\hat{\boldsymbol{p}}_{\text{KF,std.}}, \hat{\boldsymbol{p}}_{\text{KF,prop.}}$	The position estimate of a Kalman filter employing the standard ranging measurement model, and of a Kalman filter employing the proposed measurement model
$\boldsymbol{q}(t)$	The unit quaternion describing the rotation $R(t)$
$q_w(t), \tilde{\boldsymbol{q}}(t)$	The scalar and vector part of the unit quaternion
Q	The covariance matrix of the process noise
$\boldsymbol{r}$	The vector pointing from the tag antenna to the anchor antenna
$\bar{\boldsymbol{r}}$	A pseudo input point
$R(t)$	The rotation from the inertial frame to the tag frame for the entire experiment
$\text{SPGP}_m(\boldsymbol{r})$	The mean of the posterior distribution of the measured range
$\text{SPGP}_{\sigma^2}(\boldsymbol{r})$	The variance of the posterior distribution of the measured range
$\boldsymbol{v}$	The velocity of the tag expressed in the inertial frame
$\boldsymbol{w}(t)$	Continuous-time Gaussian process noise
$\boldsymbol{x}(t)$	The tag's state trajectory over the entire experiment
$\boldsymbol{x}^*(t)$	The fitted state trajectory over the entire experiment
$z_{\text{r}}, \boldsymbol{z}_{\text{acc}}, \boldsymbol{z}_{\text{gyr}}$	The measured range, the measured acceleration and the measured angular rate
$\boldsymbol{z}_{\text{r}}$	The vector of stacked measured ranges
$\mathcal{Z}_{\text{r}}, \mathcal{Z}_{\text{acc}}, \mathcal{Z}_{\text{gyr}}$	The sets of range, acceleration and angular rate measurements
$\alpha_{\text{A}}, \beta_{\text{A}}, \alpha_{\text{T}}, \beta_{\text{T}}$	The azimuth and elevation angle on the anchor and the tag
$\Gamma, \bar{\Gamma}$	The matrices of stacked range vectors and stacked pseudo input points
$\eta_{\text{r}}, \boldsymbol{\eta}_{\text{acc}}, \boldsymbol{\eta}_{\text{gyr}}$	The measurement noise in the range, in the acceleration and in the angular rate measurements
$\theta_0, \theta_1, \theta_2$	The parameters of the covariance function
Θ	The parameters of the Gaussian process
Θ^*	The fitted parameters of the Gaussian process
$\sigma_{\text{r}}, \Sigma_{\text{acc}}, \Sigma_{\text{gyr}}$	The variance of the measurement noises
$\Phi(\boldsymbol{x}(t))$	The state evolution function

VII. OUTLOOK

A framework was presented for improving real-world UWB localization accuracy by compensating for biases and other effects caused by non-idealities in the antennas and other environmental influences.

The measured range was modeled assuming only a dependence on the relative pose of two ranging modules. Other dependencies, such as a dependence on the received waveform or the received signal strength, and other kernels should be investigated as well and might lead to a further improvement in the localization accuracy, as only part of the range measurement error could be captured by the range measurement model presented. Parametric models not based on Gaussian processes should also be tested as they could lead to simpler, more accessible models.

The model obtained could also be improved by increasing the signal to noise ratio of the IMU measurements. This could be done by better damping the IMU or by flying more aggressive maneuvers during data collection.

Finally, it would also be interesting to apply the presented framework to other sensors with only partially known measurement functions and heteroscedastic noise.

APPENDIX A
NOTATION

Generally, scalars and vectors are denoted by small, standard font and small, bold font letters, respectively. Matrices are written with capitalized Latin or Greek letters. Table 3 lists the symbols in alphabetical order. The subscripts $i, j \in \mathcal{Z}_{\text{r}}$, $k \in \mathcal{Z}_{\text{acc}}$, and $l \in \mathcal{Z}_{\text{gyr}}$ denote the values of variables for specific measurements. E.g., for a specific range measurement $i \in \mathcal{Z}_{\text{r}}$, we denote the measured range as $z_{\text{r},i}$, the vector connecting the two ranging antennas as $\boldsymbol{r}_i$, and the measurement time as t_i.

APPENDIX B
VALID COVARIANCE FUNCTION

The function $k(\boldsymbol{r}, \boldsymbol{r}')$ given in (7) must be positive semi-definite to be a valid covariance function [14, p. 79]. Instead of showing the positive semi-definiteness directly, we show that the function $k(\boldsymbol{r}, \boldsymbol{r}')$ is a multiplication of two valid covariance functions, which also results in a valid covariance function [12, p. 296]. The proposed covariance function can be decomposed into two factors

$$k(\boldsymbol{r}, \boldsymbol{r}') = \theta_0 \exp\left(-\frac{1 - \frac{\boldsymbol{r}^\mathsf{T}\boldsymbol{r}'}{\|\boldsymbol{r}\|\|\boldsymbol{r}'\|}}{\theta_1} - \left(\frac{\|\boldsymbol{r}\| - \|\boldsymbol{r}'\|}{\theta_2}\right)^2\right)$$

$$= \sqrt{\theta_0} \exp\left(-\frac{1 - \frac{\boldsymbol{r}^\mathsf{T}\boldsymbol{r}'}{\|\boldsymbol{r}\|\|\boldsymbol{r}'\|}}{\theta_1}\right)$$

$$\times \sqrt{\theta_0} \exp\left(-\left(\frac{\|\boldsymbol{r}\| - \|\boldsymbol{r}'\|}{\theta_2}\right)^2\right). \quad (27)$$

While the second factor is the standard squared exponential covariance function for the inputs $\|\boldsymbol{r}\|$, $\|\boldsymbol{r}'\|$, the first factor

can be rewritten in the following form

$$\sqrt{\theta_0} \exp\left(-\frac{1 - \frac{r^\mathsf{T} r'}{\|r\|\|r'\|}}{\theta_1}\right) = c_1 \exp\left(c_2 \underbrace{\phi(r)^\mathsf{T}}_{\mathbf{x}^\mathsf{T}} \underbrace{\phi(r')}_{\mathbf{x}'}\right), \quad (28)$$

with

$$c_1 = \sqrt{\theta_0} \exp\left(\frac{-1}{\theta_1}\right) \qquad (29)$$

$$c_2 = \frac{1}{\theta_1} \qquad (30)$$

$$\phi(r) : r \to \frac{r}{\|r\|}. \qquad (31)$$

Starting from the valid linear covariance function $k(\mathbf{x}, \mathbf{x}') = \mathbf{x}^\mathsf{T}\mathbf{x}'$, and using the techniques for constructing new valid covariance functions [12, p. 296], it can be shown that also the second factor is a valid covariance function.

APPENDIX C
IMU MEASUREMENT MODELS

Given the continuous state trajectory $\mathbf{x}(t)$ of the tag with $p(t)$ the position expressed in the inertial frame, $q(t)$ the unit quaternion describing the rotation $R(t)$ from the inertial frame to the tag frame, and $b_{\mathrm{acc}}(t)$ the bias of the accelerometer, the measured acceleration of measurement $k \in \mathcal{Z}_{\mathrm{acc}}$ is

$$z_{\mathrm{acc}, k} = \underbrace{R(q(t_k))^\mathsf{T} (\ddot{p}(t_k) - g) + b_{\mathrm{acc}}(t_k)}_{h_{\mathrm{acc}}(\mathbf{x}(t_k))} + \eta_{\mathrm{acc}, k}, \quad (32)$$

where $R(q(t_k))$ denotes the rotation matrix corresponding to the rotation given by the unit quaternion $q(t_k)$, g denotes the gravitational acceleration expressed in the inertial frame and $\eta_{\mathrm{acc},k}$ is measurement noise sampled from a Gaussian distribution $\mathcal{N}(\eta_{\mathrm{acc}, k}|0, \Sigma_{\mathrm{acc}})$.

Denoting with $b_{\mathrm{gyr}}(t)$ the bias of the gyroscope, with $\tilde{q}(t)$ the vector part and with $q_w(t)$ the scalar part of the unit quaternion $q(t)$, the measured angular rate of measurement $l \in \mathcal{Z}_{\mathrm{gyr}}$ can be written as [28, eq. (322)]

$$z_{\mathrm{gyr}, l} = 2\left(q_w(t_l)\dot{\tilde{q}}(t_l) - \dot{q}_w(t_l)\tilde{q}(t_l) + \tilde{q}(t_l) \times \dot{\tilde{q}}(t_l)\right) + b_{\mathrm{gyr}}(t_l) + \eta_{\mathrm{gyr}, l}, \quad (33)$$

where $\eta_{\mathrm{gyr}}(t_l)$ denotes measurement noise from a Gaussian distribution $\mathcal{N}(\eta_{\mathrm{gyr}, l}|0, \Sigma_{\mathrm{gyr}})$.

REFERENCES

[1] A. Alarifi *et al.*, "Ultra wideband indoor positioning technologies: Analysis and recent advances," *Sensors*, vol. 16, no. 5, p. 707, May 2016.

[2] D. Porcino and W. Hirt, "Ultra-wideband radio technology: Potential and challenges ahead," *IEEE Commun. Mag.*, vol. 41, no. 7, pp. 66–74, Jul. 2003.

[3] A. Ledergerber and R. D'Andrea, "Ultra-wideband range measurement model with Gaussian processes," in *Proc. IEEE Conf. Control Technol. Appl. (CCTA)*, Aug. 2017, pp. 1929–1934.

[4] J. Khodjaev, Y. Park, and A. S. Malik, "Survey of NLOS identification and error mitigation problems in UWB-based positioning algorithms for dense environments," *Ann. Telecommun.*, vol. 65, nos. 5–6, pp. 301–311, 2010.

[5] Y. Shen and M. Z. Win, "Effect of path-overlap on localization accuracy in dense multipath environments," in *Proc. IEEE Int. Conf. Commun.*, May 2008, pp. 4197–4202.

[6] A. Dumoulin, "A study of integrated UWB antennas optimised for time domain performance," Ph.D. dissertation, Dublin Inst. Technol., Dublin, Ireland, 2012.

[7] R. Ye, S. Redfield, and H. Liu, "High-precision indoor UWB localization: Technical challenges and method," in *Proc. IEEE Int. Conf. Ultra-Wideband*, vol. 2, Sep. 2010, pp. 1–4.

[8] L. Zwirello, T. Schipper, M. Harter, and T. Zwick, "UWB localization system for indoor applications: Concept, realization and analysis," *J. Elect. Comput. Eng.*, vol. 2012, May 2012, Art. no. 849638, doi: 10.1155/2012/849638.

[9] H. Wymeersch, S. Marano, W. M. Gifford, and M. Z. Win, "A machine learning approach to ranging error mitigation for UWB localization," *IEEE Trans. Commun.*, vol. 60, no. 6, pp. 1719–1728, Jun. 2012.

[10] V. Savic, E. G. Larsson, J. Ferrer-Coll, and P. Stenumgaard, "Kernel methods for accurate UWB-based ranging with reduced complexity," *IEEE Trans. Wireless Commun.*, vol. 15, no. 3, pp. 1783–1793, Mar. 2016.

[11] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Nov. 2013, pp. 1280–1286.

[12] C. M. Bishop, *Pattern Recognition and Machine Learning* (Information Science and Statistics), 8th ed. New York, NY, USA: Springer, 2009.

[13] P. Corbalán and G. P. Picco, "Concurrent ranging in ultra-wideband radios: Experimental evidence, challenges, and opportunities," in *Proc. Int. Conf. Embedded Wireless Syst. Netw.*, Feb. 2018, pp. 55–66.

[14] C. E. Rasmussen and C. K. I. Williams, "Gaussian processes in machine learning," in *Advanced Lectures on Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.

[15] E. Snelson and Z. Ghahramani, "Sparse Gaussian processes using pseudo-inputs," in *Advances in Neural Information Processing Systems 18*. Cambridge, MA, USA: MIT Press, 2006, pp. 1257–1264.

[16] J. Quiñonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *J. Mach. Learn. Res.*, vol. 6, pp. 1939–1959, Dec. 2005.

[17] "Antenna delay calibration of DW1000-based products and systems, version 1.2," DecaWave, Appl. Note APS014. Accessed: Dec. 9, 2018. [Online]. Available: https://www.decawave.com/application-notes/

[18] P. Furgale, C. H. Tong, T. D. Barfoot, and G. Sibley, "Continuous-time batch trajectory estimation using temporal basis functions," *Int. J. Robot. Res.*, vol. 34, no. 14, pp. 1688–1710, Dec. 2015.

[19] M.-J. Kim, M.-S. Kim, and S. Y. Shin, "A general construction scheme for unit quaternion curves with simple high order derivatives," in *Proc. 22nd Annu. Conf. Comput. Graph. Interact. Techn. (SIGGRAPH)*, 1995, pp. 369–376.

[20] H. Sommer, J. R. Forbes, R. Siegwart, and P. Furgale, "Continuous-time estimation of attitude using B-splines on lie groups," *J. Guid., Control, Dyn.*, vol. 39, no. 2, pp. 242–261, 2016.

[21] M. Abadi *et al.* TensorFlow: *Large-Scale Machine Learning on Heterogeneous Distributed Systems*. Accessed: Dec. 9, 2018. [Online]. Available: http://download.tensorflow.org/paper/whitepaper2015.pdf

[22] DecaWave. *DWM 1000 Datasheed, Version 1.7*. Accessed: Dec. 9, 2018 [Online]. Available: https://www.decawave.com/product-documentation

[23] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadrocopter state estimation," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, May 2015, pp. 1730–1736.

[24] J. Ko and D. Fox, "GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models," *Auto. Robots*, vol. 27, no. 1, pp. 75–90, Jul. 2009.

[25] S. Lupashin, M. Hehn, M. W. Mueller, A. P. Schoellig, M. Sherback, and R. D'Andrea, "A platform for aerial robotics research and demonstration: The flying machine arena," *Mechatronics*, vol. 24, no. 1, pp. 41–54, 2014.

[26] A. Ledergerber, M. Hamer, and R. D'Andrea, "A robot self-localization system using one-way ultra-wideband communication," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2015, pp. 3131–3137.

[27] Z. Chen and B. Wang, "How priors of initial hyperparameters affect Gaussian process regression models," *Neurocomputing*, vol. 275, pp. 1702–1710, Jan. 2018.

[28] M. D. Shuster, "Survey of attitude representations," *J. Astron. Sci.*, vol. 41, pp. 439–517, Oct. 1993.

ANTON LEDERGERBER received the B.Sc. and M.Sc. degrees in mechanical engineering from ETH Zürich, in 2012 and 2015, respectively. During his studies, he also completed a business oriented Unitech degree and coordinated the control framework development for the unmanned blimp Skye. He is currently pursuing the Ph.D. degree with the Institute for Dynamic Systems and Control. His research is focused on estimation, ultra-wideband localization, and applications of machine learning techniques to describe physical effects which are hard to model using first principles.

RAFFAELLO D'ANDREA received the B.Sc. degree in engineering science from the University of Toronto in 1991 and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology in 1992 and 1997, respectively. He was an Assistant Professor and, then an Associate Professor, at Cornell University from 1997 to 2007. While on leave from Cornell University, from 2003 to 2007, he co-founded Kiva Systems (now Amazon Robotics), where he led the systems architecture, robot design, robot navigation and coordination, and control algorithms development. He is currently a Professor of dynamic systems and control at ETH Zürich and also the Chairman of the Board at Verity Studios AG.

• • •