

Semi-Unsupervised Mitigation of Human Body Shadowing for Indoor UWB pedestrian tracking

Cedric De Cock
Dept. of Information Technology
imec-WAVES, Ghent University
Ghent, Belgium
cedric.decock@ugent.be

Emmeric Tanghe
Dept. of Information Technology
imec-WAVES/Ghent University
Ghent, Belgium
emmeric.tanghe@ugent.be

Wout Joseph
Dept. of Information Technology
imec-WAVES/Ghent University
Ghent, Belgium
wout.joseph@ugent.be

David Plets
Dept. of Information Technology
imec-WAVES/Ghent University
Ghent, Belgium
david.plets@ugent.be

Abstract—in Ultra Wideband (UWB), large ranging errors occur under Non-Line-of-Sight (NLoS) conditions, which significantly degrades positioning accuracy. Human body shadowing (HBS) is a specific case of NLoS, which is a prominent error source for on-body UWB positioning. This work presents a tracking algorithm based on a Particle Filter (PF), designed to mitigate HBS-induced positioning errors by using an orientation-adaptive measurement model, consisting of a bank of Gaussian Mixture Models. The relative orientation is derived from Inertial Measurement Unit (IMU) data, and predicted positions from the tracking algorithm itself. We propose a second tracking algorithm in order to train the adaptive measurement model in a semi-supervised way, eliminating the need for accurate ground truth. The proposed algorithm outperforms a state of the art algorithm by an average of 11% (unsupervised) to 39% (supervised) in an experimental evaluation.

Index Terms—Gaussian Mixture Model, human body shadowing, indoor positioning, IMU, NLoS, unsupervised learning, UWB

I. INTRODUCTION

For the last two decades, a large amount of research has been directed towards detection and mitigation of non-line-of-sight (NLoS) effects in Ultra-Wideband (UWB) ranging, in which the direct path between tag and anchor is obstructed by a wall, the human body, or other object. In NLoS conditions, it is possible that a diffracted or reflected path component is seen as the first path component, which adds a bias to the estimated range [1]. Consequently, these NLoS-induced the performance range errors degrade range-based positioning.

Human Body Shadowing (HBS) is the specific NLoS case in which a human body obstructs the line-of-sight (LoS) path. Several works have already investigated HBS, and have identified it as an important cause of NLoS-induced range errors in on-body UWB ranging, i.e. when the tag is attached to a person's body [2], [3]. Regardless of the specific setup, a common pattern in the range error distribution has been observed by several works [3], [4]. First, there are no significant errors when there is no direct obstruction, i.e. when the tag is in between the anchor and body. As the person

is turning, the direct part gradually becomes obstructed, and the measured range becomes biased up to 10 – 30 cm, due to the first path component diffracting around the body. These diffracting or creeping waves are heavily mitigated, thus when the body is fully in between the tag and anchor, the creeping waves remain often undetected [2]. In that case, multipath components who have reflected off other surfaces are detected instead, causing outliers up to several meters, which heavily skews the otherwise Gaussian range error distribution.

Despite the mentioned efforts of characterizing the HBS effects, and the large amount of works on NLoS detection and mitigation in UWB context, there is limited work on exploiting the knowledge of this orientation-dependent effect to improve on-body UWB positioning accuracy. In [5], we proposed a robust HBS mitigation algorithm for static positioning, assisted by an IMU. While [5] focused on static positioning, this work is aimed at dynamic on-body UWB-Two-Way Ranging (TWR) positioning, for which a new positioning algorithm is designed and implemented. The general approach of both methods is the same, as the Inertial Measurement Unit (IMU) is used in order to estimate the relative orientation of the body with respect to the anchor and tag, after which a more accurate position is estimated using knowledge of the HBS effect on the range measurements. To mitigate the HBS effect, we present a measurement model based on a bank of Gaussian Mixture Models (GMMs), which characterizes the UWB range errors as a function of the IMU-based orientation. However, training these GMMs requires accurate ground truth data matching the orientation to the ranging error, which can be expensive to obtain. Therefore, in line with our philosophy in [5] of easy deployability, we propose the usage of a second algorithm, which can train the GMM-based measurement model, i.e. the UWB range error model, of the first system without the need for accurate ground truth. Together, this forms a semi-supervised on-body UWB positioning system, which mitigates HBS effects by characterizing the orientation-dependent range error distribution almost independently.

II. METHOD

This section describes the used hardware, the environment in which the measurements were performed, followed by the estimation of the relative orientation and a system overview.

A. Measurement configuration

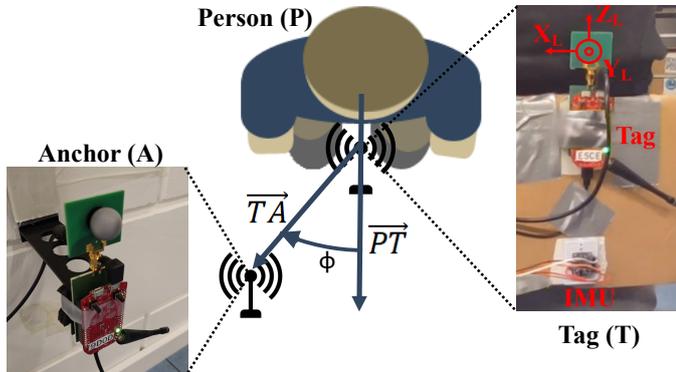


Fig. 1: Visualization of vectors related to estimation of the relative tag orientation ϕ , with images of the hardware used during the measurement campaign.

1) *Hardware and data collection*: The experiment setup is visualized in Figure 1, which depicts a person wearing an UWB tag at the belly. A picture of this tag is provided on the right side of Figure 1, showing an attached IMU, while the picture on the left shows an UWB anchor. The latter is identical in terms of hardware, and can be seen with an Infrared (IR) marker on its antenna for accurate calibration of the anchor position. The UWB hardware used is the Wi-Pos platform [6]. Wi-Pos consists of a Decawave transceiver for UWB ranging, and a Zolertia Re-mote with sub-GHz radio for orchestration of the TWR scheme and for reporting the estimated range, which is performed at a rate of 23 Hz. The IMU used is the Adafruit BNO055, which consists of a 3-axis accelerometer, gyroscope, and magnetometer, as well as a built-in Attitude and Heading Reference System (AHRS) module, which fuses the inertial measurements to obtain accurate orientation at 100 Hz. The UWB and IMU data are read by a Raspberry Pi, which immediately sends them a local MQTT broker over Wi-Fi. Furthermore, accurate ground truth is collected by a Qualisys motion capture (mocap) system at 90 Hz, which is also sent to the MQTT broker. This ground truth consist of accurate positions as well as orientation of the tracked object. Figure 1 shows in red how the local reference frame of the mocap-tracked body is defined, having the tag antenna as the local origin, and the local Y-axis coinciding with \vec{PT} .

2) *Environment and trajectories*: The experiments are performed in a 11 m x 9 m rectangular open area in the Internet of Things (IoT) of the IDLab research group, featuring eight Wi-Pos UWB anchors. Two UWB anchors are vertically stacked at each of the corners, as shown by the red dots in Figure 2. The axes of Figure 2 show how the global coordinate system is defined, in which the Y_G -axis coincides with North. A training

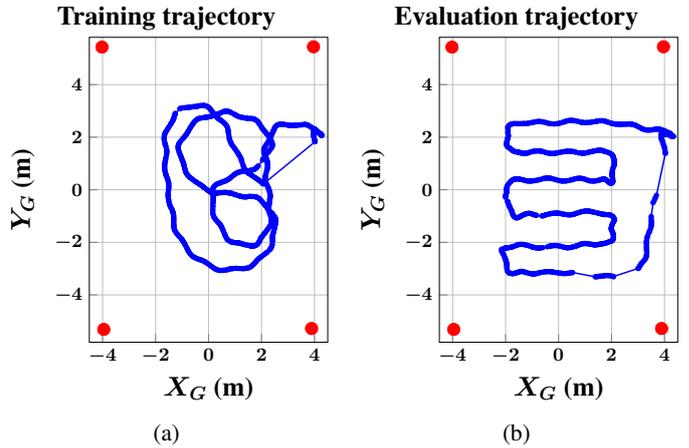


Fig. 2: Ground truth scatterplots of the training (a) and evaluation (b) trajectories. Straight lines denote parts of the trajectories with missing ground truth. UWB anchors are marked in red, with each marker representing two anchors.

and evaluation trajectory were designed, of which the mocap-provided ground truth are shown on the blue scatterplots in Figure 2a and Figure 2b respectively. Each trajectory is travelled five times, each with a duration of 50 – 60 s, or an average of 1160 UWB range measurements. The pedestrian carried the tag at waist height, and walked casually along the predefined trajectory, with some slight variations during each repetition, e.g. sometimes taking a turn wider or sharper. These variations do not require any extra work, since the mocap system records the ground truth position with mm-level accuracy while the UWB range measurements are taking place.

B. Orientation estimation

Figure 1 illustrates the discussed relative orientation for our setup. The relative orientation is defined as the smallest angle $\phi \in [0^\circ, 180^\circ]$ between the 2D vectors \vec{TA} and \vec{PT} in the global XY-plane, in which T is the tag position, A is the anchor position, and P represents the person's center of mass. Given the range of possible ϕ values, it is assumed that the orientation-dependent range error distribution is symmetrical with respect to \vec{PT} . To comply with this assumption, both the tag and IMU are placed at a fixed position at the center of the belly, and the antenna surface is placed perpendicular to the \vec{PT} . Furthermore, \vec{TA} can be estimated given the anchor position and a predicted tag position, whereas \vec{PT} can be derived from the yaw angle estimated by the IMU's AHRS. For more details, we refer to our prior work [5].

C. System design

1) *Particle Filtering*: The design of the proposed positioning system is explained using the flowgraph in Figure 3. The vertical dashed line divides the system into an offline phase (left) for training the measurement noise model of the online phase (right). As mentioned in Section I, each of these two parts is a separate tracking algorithm, of which the online one is based on a Particle Filter (PF), while its

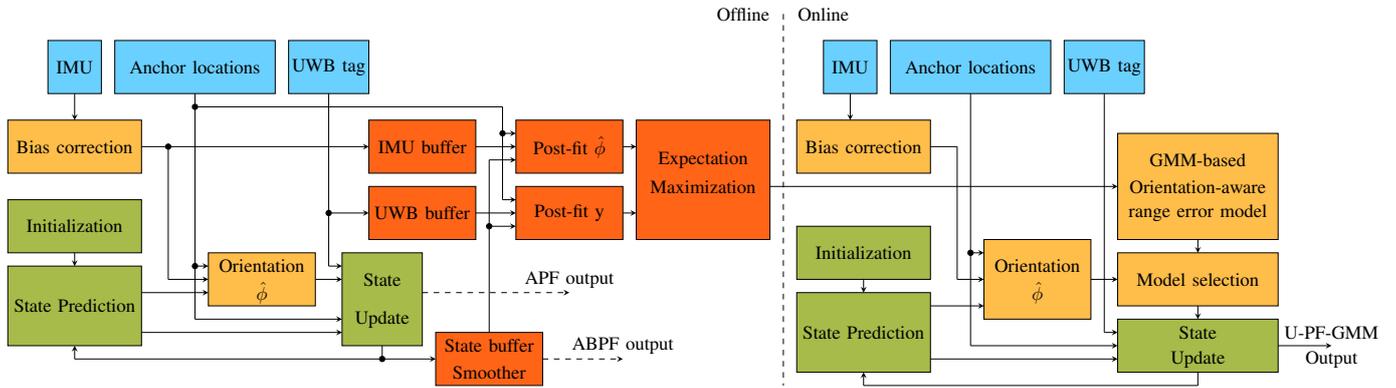


Fig. 3: Flowchart of the Human Body Shadowing mitigation approach. The blue blocks represent the data input to the system, whereas the green blocks represent the typical steps of the tracking algorithm. The red blocks represent the unsupervised training of the range error model, which is part of the offline phase left of the vertical dashed line. The yellow blocks represent the orientation estimation and usage of the proposed orientation-adaptive measurement model.

offline counterpart is a smoother, more specifically based on a Backtracking Particle Filter (BPF) [7]. While it can produce real-time state estimations identical to a regular PF, the BPF remembers each particle's predecessor, which allows it to backtrack through previous particle generations and re-estimate a previous position based only on the particles with surviving offspring. Both algorithms consist of the same three elements, represented by the green blocks in Figure 3, starting at the initialization. In this work, each algorithm waits until one range has been received from four anchors. Then, the initial position is calculated using Linearized Least Squares (LLS) multilateration [8], after which a set of particles is sampled from a Gaussian distribution centered around this LLS position. Aside from a position, each particle's state also includes a velocity, which is initially set to $0 \frac{m}{s}$ as it is assumed that the experiment starts from a stationary pose. After initialization, a two-step cycle is performed for each incoming UWB range measurement. The prediction step adds artificial noise to the velocity, and predicts a new position for each particle using a Constant Velocity (CV) process model. The update step calculates the (pre-fit) range residual for each particle, i.e. the difference between the measured range and the distance between the anchor and the particle. Each particle is then weighted by plugging its range residual into a measurement likelihood function, after which a weighted position is calculated.

For best performance of the filter, the likelihood function should resemble the Probability Density Function (PDF) of the measurement noise. A Gaussian PDF is often used as likelihood function, and works well for UWB positioning in LoS conditions. However a Gaussian likelihood function is unsuitable for use in NLoS conditions, in which the range error PDF is heavily skewed. On top of that, in HBS conditions, the range error PDF is dependent on the orientation of the user. Therefore, the PF can benefit from an orientation-adaptive measurement model.

2) *Orientation-adaptive measurement model*: The difference between the offline and online algorithms lies in the likelihood function, which is implied by the yellow blocks in Figure 3. In both algorithms, the orientation ϕ is estimated for each particle, based on each particle's position, the anchor position, and the yaw angle provided by the IMU (after a bias correction). For the online PF specifically, the ϕ value of each particle is used to select a GMM as likelihood function, which resembles the range error distribution for that value of ϕ . For low ϕ values, the GMM's PDF is Gaussian-like, while the PDF is heavily skewed for $\phi \rightarrow 180^\circ$. This online PF algorithm with orientation-adaptive measurement model based on GMMs, is denoted as PF-GMM when the measurement noise model is trained on mocap ground truth (supervised, benchmark), and is denoted as U-PF-GMM when it is trained on output of the offline algorithm (unsupervised). The offline algorithm evidently does not have this GMM-based measurement model, but uses an adaptive Gaussian PDF, of which the variance is doubled when $\hat{\phi}$ surpasses a threshold ϕ_c . The latter is inspired by [9], in which the measurement noise variance of an Extended Kalman Filter (EKF) is increased by thresholding the skewness of the Channel Impulse Response. Furthermore, since the Adaptive Backtracking Particle Filter (ABPF) can provide real-time state estimations, Figure 3 shows an intermediate output called Adaptive Particle Filter (APF) flowing out of the ABPF's update step.

3) *Semi-supervised training*: A regular BPF requires the tuning of two parameters: the noise variance of the process model and of the measurement model, to which the ABPF adds the threshold ϕ_c . Our system is semi-supervised, because these parameters still have to be tuned manually. However, unlike the PF-GMM with its GMM-based measurement model, the two variances of the BPF can be tuned without the need for accurate ground truth, by analyzing the BPF's range residuals, visually validating the convergence, etc. Furthermore, ϕ_c is a robust parameter, as good (and similar) performance is achieved for any value of $\phi_c \in [110^\circ, 160^\circ]$.

The red blocks in Figure 3 denote the (semi-)unsupervised training of the GMM-based measurement model. While the ABPF processes the training data, all estimated states are kept in memory, i.e. in the state buffer, until all predict-update cycles are finished. As mentioned, all positions are then re-estimated, which makes the ABPF a smoother algorithm. A second intermediate output is shown at the smoother block in Figure 3. This output, along with the stored IMU and UWB measurements, serve as the input for training the measurement model. More specifically, a post-fit range residual y' and orientation $\hat{\phi}'$ are calculated for each smoothed position, which results in a training dataset consisting of pairs of y' and $\hat{\phi}'$ values. It is expected that these post-fit parameters are closer to their ground truth counterparts than the pre-fit values, which are calculated before each update step.

Finally, to obtain a bank of GMMs, a random subset of range residuals is sampled from the training data for $\phi \in [0^\circ, 180^\circ]$. For a given ϕ , the range residuals are weighted based on their corresponding $\hat{\phi}'$ value, in such a way that residuals with a $\hat{\phi}'$ close to ϕ have a higher chance of being selected. For each subset, several GMMs with a differing amount of Gaussian components are then fitted using the Expectation-Maximization (EM) algorithm as implemented in [10]. The GMM with the lowest Bayesian Information Criterion for a particular subset is then selected for the corresponding ϕ value, resulting in a bank of 181 GMMs.

4) *Evaluated algorithms:* Several algorithms have been described in this section, which are evaluated in Section III. The main evaluated algorithm is the U-PG-GMM, the proposed PF algorithm which uses the orientation-adaptive GMM-based measurement model trained by the ABPF. The ABPF, which uses a simple threshold ϕ_c to increase the measurement noise variance of its Gaussian measurement model, is evaluated too. The ABPF can produce real-time results by simply not performing the backtracking/smoothing step, in which case it operates as a filter. This configuration is denoted as the APF, of which the performance is also evaluated. The APF is the most important benchmark for the U-PF-GMM, since outperforming the APF means that the unsupervised training with the ABPF is justified. The PF-GMM is used to evaluate the highest possible performance of a PF with the GMM-based measurement model, by training this model on mocap-provided ground truth range errors and ϕ values. Furthermore, a state-of-the-art algorithm PF-Ref [4] is implemented. The PF-GMM uses a ϕ threshold to switch between a Gaussian with adaptive mean, and Gamma PDF. More specifically, a Linear Regression model is used to predict the Gaussian mean as a function of $\hat{\phi}$. The PF-Ref's measurement model is therefore also trained using the mocap ground truth of the training trajectory. Finally, a traditional PF is used as a benchmark for all other orientation-adaptive algorithms. This PF is also configured as a smoother, i.e. BPF, in order to evaluate the advantage of using the ABPF for training the U-PF-GMM's measurement model. A short description of the discussed tracking algorithms and their prominent differences in the context of this work, are summarized in Table I.

III. RESULTS

The results presented in this section are based on all five recordings of the evaluation trajectory. Also, since PFs are non-deterministic, each PF is run ten times per trajectory using 400 particles. A position error is defined as the Euclidean distance in the XY-plane between the position estimated by a specified tracking algorithm for a given UWB measurement, and the mocap-provided position of which the timestamp lies closest to the timestamp of that UWB measurement. The position error is left out of the results if the closest mocap timestamp is more than 0.1 s away from the UWB measurement's timestamp, which occasionally happens when the mocap system has failed to track the tag.

A. Evaluation of unsupervised range error characterization

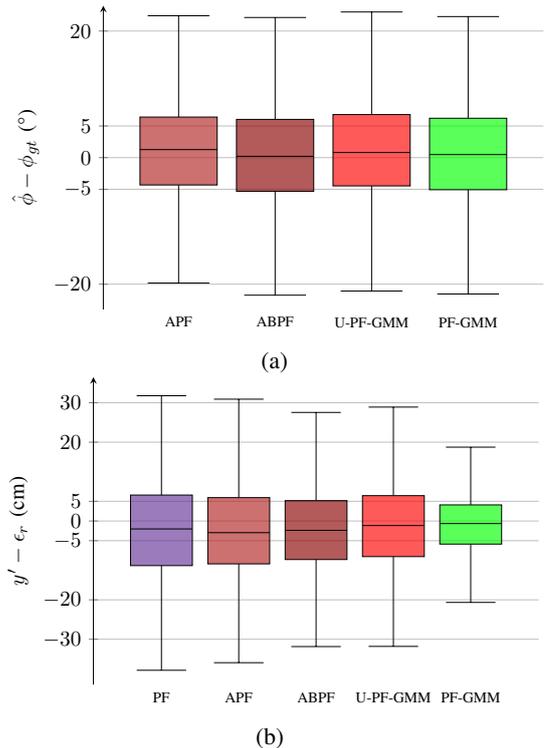


Fig. 4: Boxplots of post-fit $\hat{\phi}$ errors (a), and range residual y' errors (b). Post-fit errors are calculated after smoothing (ABPF), or after the state update (others).

Figure 4a shows boxplots of the post-fit $\hat{\phi}$ errors for the ABPF, U-PF-GMM and PF-GMM. The post-fit $\hat{\phi}$ errors are calculated by subtracting the mocap-based ϕ_{gt} from $\hat{\phi}$, in which the latter is estimated after the filter algorithm has executed the state update. In the case of the smoother algorithm (ABPF), the post-fit $\hat{\phi}$ error is calculated after the smoothed states have been acquired. From Figure 4a, it appears the $\hat{\phi}$ error is not much affected by the tracking algorithm. This is because the distance of any point of both trajectories to the closest anchor is larger than 2 m, as shown in Figure 2. Therefore, even if the position performance differs between the type of tracking algorithm, $\hat{\phi}$ is not much affected, thus

TABLE I

Algorithm	measurement model	accurate ground	Description
PF	Gaussian	no	standard Particle Filter algorithm with constant velocity process model
BPF	Gaussian	no	smoother variant of PF
APF	Adaptive Gaussian	no	PF with orientation-adaptive measurement model by thresholding $\hat{\phi}$
ABPF	Adaptive Gaussian	no	smoother variant of APF
PF-Ref [4]	Gaussian/Gamma	yes	state-of-the-art PF algorithm with orientation-adaptive measurement model
PF-GMM	GMM bank	yes	PF with orientation-adaptive measurement model using a bank of GMMs
U-PF-GMM	GMM bank	no	Unsupervised PF-GMM trained without ground truth by output of ABPF

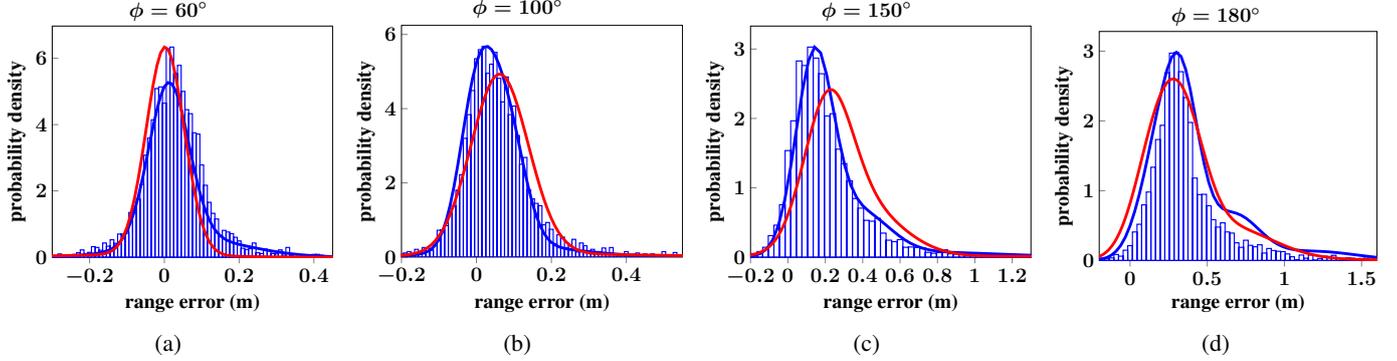


Fig. 5: Histograms of UWB range error subsets with GMM-based PDFs fitted on mocap data (blue), and on subsets of range residuals provided by a smoother algorithm (red). Each subset is sampled around ϕ_{gt} (blue) or $\hat{\phi}$ (red).

depends mostly on the IMU. This is also complemented by the fact that the discussed tracking algorithms generally do not produce enough outliers (in this experiment) that can significantly affect $\hat{\phi}$ (see Section III-B). Regardless, the IMU enables accurate estimation of ϕ , achieving a median absolute error under 6° for each algorithm. This can still cause problems for a threshold-based adaptive measurement model, such as that of the A(B)PF or the PF-Ref [4], but is less of a problem for the PF-GMM. Furthermore, the $\hat{\phi}$ as estimated by the ABPF is unbiased unlike the APF, which already justifies the use of the smoother variant for training the GMM-based measurement model.

Figure 4b shows the post-fit range residual errors of the same filter algorithms as Figure 4a, as well as the standard PF. These residual errors are calculated by subtracting the range error ϵ_r from the post-fit range residual y' , i.e. the range residual calculated after the state update step, or after smoothing if applicable. Figure 4b makes clear the advantage of the APF over the PF algorithm, and of the ABPF over the APF for training the GMM-based model. Especially the ABPF clearly achieves a smaller error for the range residual, allowing the training of a measurement model which resembles the ground truth trained model of the PF-GMM more closely. Note that the range residual is consistently underestimated by all algorithms, This is normal, because the UWB range is typically overestimated under NLoS conditions. Consequently, the tracking algorithm incorporating a range measurement is more likely to estimate a position further away from the anchor than the actual position, resulting in an underestimated post-fit range residual. Training the GMM-based measurement noise

model on biased training data is not desirable. However, this problem is partly solved by assuming the range error resembles a zero-mean Gaussian PDF for $\phi < 90^\circ$. The PDFs estimated by the EM algorithm can then be shifted by the opposite of the y' value at the PDF's peak. Furthermore, the PF-GMM achieves the smallest underestimation of y' in absolute terms (1 cm), and also has the smallest error standard deviation ($\sigma_{y'} = 8.8$ cm), implying that the GMM-based model is a good representation of the true range error distribution, and that the PF-GMM achieves high positioning accuracy.

Figure 5 shows the histogram of UWB range errors from the training dataset sampled around four ϕ values based on their corresponding ϕ_{gt} as described in Section II. The blue curves represent the PDFs of the GMMs fitted on these range errors by the EM algorithm, while the red curves show the PDFs of GMM trained on subsets of range residuals y' of ABPF. The subsets of y' are in turn sampled around ϕ values based on their corresponding $\hat{\phi}$ value. Both types of GMMs conform to the general pattern of orientation-dependent range error distributions in HBS context discussed in Section I, in which a zero-mean Gaussian-like PDF at $\phi = 0^\circ$ transitions to a positively shifted, heavily skewed PDF for $\phi \rightarrow 180^\circ$. For low and high values of ϕ , the models trained on the ABPF output show a good overlap with the models trained on ground truth data. However, this overlap is achieved to a lesser extent for medium ϕ values. The reason is that when the person is not maneuvering, the tracking algorithm can estimate the general direction of the person relatively well thanks to the CV process model. However due to the range measurements being biased, the estimated trajectory tends to have an offset to the side of

the true trajectory. In this particular measurement setup, the tag is carried on the belly, and the person neither walked sideways nor backwards. Therefore, the position errors to the side of the true trajectory affect the range residual y' less when the anchor is located in the extension of the trajectory, i.e. close to either $\phi = 0^\circ$ or $\phi = 180^\circ$.

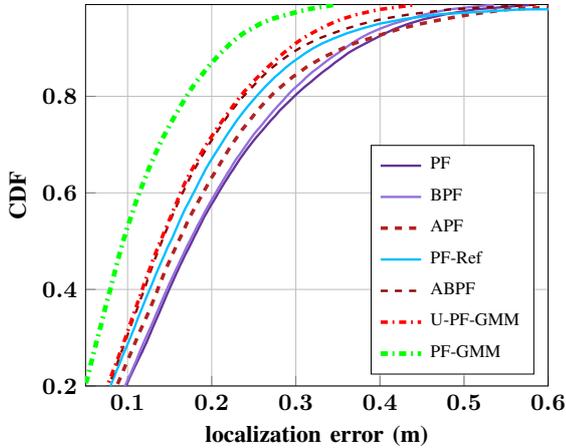


Fig. 6: CDFs of localization errors of discussed algorithms

B. Evaluation of proposed orientation-based positioning algorithm

Figure 6 shows the position error Cumulative Distribution Functions (CDFs) of the discussed tracking algorithms. The proposed PF-GMM and U-PF-GMM algorithms are compared to the well-known PF algorithm, to a state-of-the-art HBS mitigation algorithm PF-Ref [4], and to the APF of which the smoother variant ABPF is used to train the U-PF-GMM's measurement model. The PF is used as the main benchmark algorithm, and achieves a median, p75, p90 and p99 error of 18 cm, 27 cm, and 37 cm respectively. For the remaining algorithms, reductions of position errors are expressed relative to the standard PF, unless specified otherwise. Its smoother variant (BPF) shows only a slight improvement, reducing the p99 to 54 cm (-4%) and reducing other percentiles by 1 cm. The APF reduces the p75 and p90 error only slightly to 24 cm (-7%) and 36 cm (-3%) respectively, while suffering from higher outliers as the p99 error increases to 61 cm ($+5\%$). However, these outliers are eliminated in the smoother variant (ABPF), and achieves a p50, p75 and p90 error of 14 cm (-22%), 21 cm (-22%), and 30 cm (-19%) respectively, outperforming the APF and smoother variant of the PF on all metrics. PF-GMM is the best performing algorithm, which manages to reduce the median, p75 and p90 error to 9 cm (-50%), 15 cm (-44%) and 22 cm (-43%) respectively. U-PF-GMM achieves a median, p75, p90, and p99 error of 14 cm, 21 cm, 29 cm, and 45 cm respectively, which is an improvement of 22% across all percentiles. Therefore, U-PF-GMM does not perform as good as PF-GMM, but outperforms all other algorithms while not needing accurate ground truth to train its measurement model, unlike the PF-GMM. Furthermore, while U-PF-GMM is only slightly better than the ABPF

by which it was trained, the latter being a smoother algorithm, does not produce real-time results. Lastly, the referenced PF (PF-Ref), of which the model is trained on our ground truth data, is outperformed by both our proposed (U-)PF-GMM algorithms. When using the measurement model parameters fitted for the experiments in [4], the PF-Ref diverges.

To investigate how much training data is needed to train U-PF-GMM's model, we split the ABPF's output in splits of varying sizes. This is done by letting the ABPF process all recorded measurements of the training trajectory, after which e.g. the first 100 post-fit range residuals y' and corresponding $\hat{\phi}$ values are taken. This simulates scenarios in which the person walked for a shorter amount of time, giving the system only a fraction of the total amount of measurements to characterize the range errors. Figure 7 shows different percentile positioning errors of U-PF-GMM as a function of the relative size of the training dataset. It is observed that the accuracy increases the most for the first 900 range measurements, after which significantly more data is needed in for marginal improvements. Lastly, it is found that a training dataset of 1500 measurements, i.e. 66 s is enough for U-PF-GMM to outperform the APF algorithm.

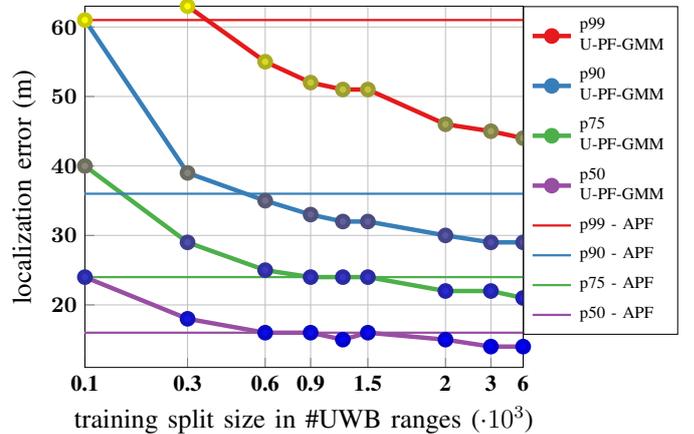


Fig. 7: localization error statistics of the U-PF-GMM as a function of the size of the unsupervised training dataset

IV. CONCLUSION AND FUTURE WORK

HBS is an important form of NLoS in on-body UWB positioning. Although the orientation dependence of the range error distribution has been observed in several works, only a limited amount of works have used this knowledge to mitigate the effects of HBS on UWB positioning. In this work, a PF algorithm is presented which employs an orientation-aware measurement model based on a bank of GMMs. This filter algorithm estimates the orientation of the UWB tag relative to the body and anchors after predicting a candidate position, after which an orientation-specific GMM is selected for the state update step. Furthermore, a semi-supervised system is proposed, in which a smoother algorithm is used to train the GMM-based model, with the goal of eliminating the need for accurate ground truth. The proposed Unsupervised PF-GMM outperforms the benchmark algorithms, achieving a

19% reduction of the p90 error after training on a 250 s recording of UWB and IMU data, and a 39% reduction when trained on ground truth data.

For future work, the authors would like to test the proposed method in a more realistic environment, which includes external NLoS sources, such as walls and metallic racks.

ACKNOWLEDGMENTS

This work is supported by the imec co-financed project UWB-IR AAA and by the Research Foundation Flanders (FWO) project under Grant G098020N. The authors would like to thank B. Van Herbruggen, N. Macoir, J. Vanhie-Van Gerwen, and E. De Poorter from IMEC-IDLab/Ghent University for their assistance during the measurement campaign.

REFERENCES

- [1] Kegen Yu, Kai Wen, Yingbing Li, Shuai Zhang, and Kefei Zhang. A novel NLOS mitigation algorithm for UWB localization in harsh indoor environments. *IEEE Transactions on Vehicular Technology*, 68(1):686–699, 2019.
- [2] Timo Kumpulainen, Matti Hamalainen, Juha Pekka Makela, and Jari Iinatti. Path loss modeling for UWB creeping waves around human body. *International Symposium on Medical Information and Communication Technology, ISMICT*, pages 44–48, 2017.
- [3] Timothy Otim, Alfonso Bahillo, Luis Enrique Díez, Peio Lopez-Iturri, and Francisco Falcone. FDTD and Empirical Exploration of Human Body and UWB Radiation Interaction on TOF Ranging. *IEEE Antennas and Wireless Propagation Letters*, 18(6):1119–1123, 2019.
- [4] Timothy Otim, Alfonso Bahillo, Luis Enrique Díez, Peio Lopez-Iturri, and Francisco Falcone. Towards sub-meter level uwb indoor localization using body wearable sensors. *IEEE Access*, 8:178886–178899, 2020.
- [5] Cedric De Cock, Sander Coene, Ben van Herbruggen, Luc Martens, Wout Joseph, and David Plets. Imu-aided detection and mitigation of human body shadowing for uwb positioning. In *2022 IEEE 12th International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–8, 2022.
- [6] Ben Van Herbruggen, Bart Jooris, Jen Rossey, Matteo Ridolfi, Nicola Macoir, Quinten Van Den Brande, Sam Lemey, and Eli De Poorter. Wi-pos: A low-cost, open source ultra-wideband (UWB) hardware platform with long range sub-GHZ backbone. *Sensors (Switzerland)*, 19(7):1–16, 2019.
- [7] widyawan, Martin Klepal, and Stéphane Beauregard. A Novel Backtracking Particle Filter for Pattern Matching Indoor Localization. In *Proceedings of the First ACM International Workshop on Mobile Entity Localization and Tracking in GPS-Less Environments, MELT '08*, pages 79–84, New York, NY, USA, 2008. Association for Computing Machinery. event-place: San Francisco, California, USA.
- [8] Abdelmoumen Norrdine. An algebraic solution to the multilateration problem. In *2012 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 11 2012.
- [9] Enrique Garcia, Pablo Poudereux, Alvaro Hernandez, Jesus Urena, and David Gualda. A robust UWB indoor positioning system for highly complex environments. *Proceedings of the IEEE International Conference on Industrial Technology*, 2015-June(June):3386–3391, jun 2015.
- [10] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.