Advancing Smartphone-based Indoor Positioning through Particle Distribution Optimization

Toni Fetzer
cronn GmbH
Bonn, Germany
toni.fetzer@cronn.de

Markus Bullmann, Steffen Kastner, and Frank Deinzer
Center for Artificial Intelligence and Robotics
Technical University of Applied Sciences Würzburg-Schweinfurt
Würzburg, Germany
{firstname.lastname}@thws.de

Marcin Grzegorzek
Institute of Medical Informatics
University of Lübeck
Lübeck, Germany
grzegorzek@imi.uni-luebeck.de

***Abstract*—Smartphone-based indoor positioning and navigation remains a challenging task, as specialized technologies such as ultra-wideband (UWB) or Wi-Fi fine-time measurement are still niche and supported by only a few flagship smartphones. Therefore, standard technologies based on RSSI measurements, mainly Bluetooth Low Energy (BLE) and Wi-Fi, are used to obtain absolute positioning information of pedestrians inside buildings. Sensor fusion methods combine this with relative information from modeling human movement using sensor data provided by the smartphone's IMU. It is also common practice to restrict this movement to the actual accessible areas of the building (e.g. restricting moving through walls), using spatial models based on the building's floor plan. Without further assumptions, this complexity inevitably leads to a non-linear and non-Gaussian state space model. A common tool for (position) estimation in such scenarios is the broad class of particle filters. However, the use of such spatial constraints accelerates the well-known problem of sample impoverishment, which in the worst case can lead to the particle filter completely losing track, getting stuck and never recovering. This work begins with a brief presentation of an award-winning Indoor Positioning System (IPS) derived from previous work. Based on this, we present several approaches using Particle Distribution Optimization (PDO) that attempt to solve the impoverishment problem and ultimately lead to better overall positioning results. In the experiments, we compare them in two different buildings under realistic conditions and discuss the results in detail.**

***Index Terms*—Particle Filter, Sensor Fusion, Indoor Positioning, Particle Distribution Optimization, Target Tracking**

I. Introduction

In recent years, many approaches addressing the problem of finding a pedestrians positions inside buildings were presented [1]–[4]. Here, smartphones are an ideal source of information to make indoor positioning technology accessible to as many people as possible. They are equipped with a variety of sensors, in particular IMU (accelerometer, gyroscope, magnetometer), barometer, Bluetooth Low Energy (BLE) and Wi-Fi. The combination of all available sensors, also known as sensor fusion, enables particularly robust and accurate position estimation. In addition, spatial models provide information about walls, stairs, doors, or other obstacles in the building [5]. Combined with the smartphone's IMU, human motion can be realistically modeled, for example, to prevent walking through walls. All this leads to a significant complexity of the dynamic system to be modeled. From a mathematical point of view, this results in a non-linear, non-Gaussian distributed state space representation in which a (hidden) state, i.e. the position of the pedestrian, has to be determined from incomplete and noisy sensor measurements over time. Since an analytical solution is not feasible, Indoor Positioning Systems (IPS) are often based on the concept of particle filters [6]–[8].

The effectiveness of IPS using particle filters is demonstrated by their performance in past competitions for smartphone-based solutions organized by the EvAAL group at various locations around the world [9]. These events emphasize realistic conditions, requiring the use of commercial devices in unprepared multi-story buildings. The best results so far have been achieved with our IPS based on a particle filter using IMU, barometer and Wi-Fi / BLE RSSI measurements [10]–[12]. Similar approaches have dominated the field in recent years. Some other notable competitors are the IPS of [13], [14] and [15].

However, most, if not all, of these systems suffer from the severe problem of sample impoverishment. Basically, the particle filter uses a set of weighted random samples (particles) to approximate a probability distribution describing the dynamical system. It works in three steps: First, new particles are drawn according to an importance distribution, most often represented by the state transition. Then, these particles are weighted by an incremental importance weight distribution based on the sensor measurements. Finally, a resampling step is applied to avoid that only a small number of particles have a significant weight and all others have negligible small weights instead. Here, a new set of equally weighted particles is drawn by duplicating high weighted particles and dropping low weighted ones. However, this leads to decreasing diversity and high concentration of particles after a resampling step, and thus to a poor approximation of the underlying probability distribution. This effect is then known as sample impoverishment.

As we have shown in [16], another major effect besides resampling that exacerbates the impoverishment problem are restrictive motion models. Imagine a pedestrian leaving a room and turning to the right into a hallway, but due to uncertain measurements, most of the particles remain in the room. Since movement through walls is restricted, the particles will be trapped inside the room, resulting in a poor approximation

of the true position, and in the worst case, all particles will be trapped, causing the filter to get stuck.

In a general sense, solving for effects that affect the approximation behavior of the particle set is known as particle distribution optimization (PDO). The purpose of this work is to present known and novel approaches that attempt to reduce or even solve the problem of sample impoverishment. Based on our award-winning IPS, we evaluate and compare them in two different buildings under realistic conditions and discuss the results in detail.

II. Indoor Positioning System

In a general sense, particle filter are a direct application of sequential Monte Carlo (SMC) to a sequence of target distribution $\tau(\boldsymbol{q}_{1:t}) = p(\boldsymbol{q}_{1:t} \mid \boldsymbol{o}_{1:t})$ and its marginals. A common way to formalize this is to use the analytical form of a recursive Bayes filter, as follows:

$$p(\boldsymbol{q}_t \mid \boldsymbol{o}_{1:t}) \propto \underbrace{p(\boldsymbol{o}_t \mid \boldsymbol{q}_t)}_{\text{evaluation}} \int \underbrace{p(\boldsymbol{q}_t \mid \boldsymbol{q}_{t-1})}_{\text{transition}} \underbrace{p(\boldsymbol{q}_{t-1} \mid \boldsymbol{o}_{1:t-1}) d\boldsymbol{q}_{t-1}}_{\text{recursion}} \quad . \tag{1}$$

It claims that given all observations $\boldsymbol{o}_{1:t}$ up to the current time t, we can compute an estimate of the Markov hidden state $\boldsymbol{q}_t$ using the principles of Bayesian inference [17]. Consequently, the particle filter uses a finite number of particles $\Upsilon_t = \{W^i_{1:t}, \boldsymbol{X}^i_{1:t}\}^N_{i=1}$ to approximate the marginal posterior $p(\boldsymbol{q}_t \mid \boldsymbol{o}_{1:t})$, where each particle is a weighted representation of a possible system state.

$$\boldsymbol{q}_t = (x, y, z, \theta), \;\; x, y, z, \theta \in \mathbb{R} \;\; , \tag{2}$$

where x, y, z represent the position in three-dimensional Euclidean space and θ the absolute heading of the user.

In short, the evaluation describes the uncertainty of the sensors used and thus provides the weight W_t of the corresponding sample $\boldsymbol{X}_t$. Assuming statistical independence between the respective sensors and their models, we define the probability density of the evaluation as

$$\begin{aligned} W^i_t \propto p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t) = \\ p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{wifi}} \;\; p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{ble}} \\ p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{baro}} \;\; p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{act}} \;\; . \end{aligned} \tag{3}$$

Here, each model has a probabilistic representation answering the same question "How likely is it that I am at this position?". The specific implementation of the sensor models is described in previous work: The Wi-Fi model $p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{wifi}}$ in [18] and [11], the BLE model $p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{ble}}$ in [8], the barometric model $p(\boldsymbol{o}_t \mid \boldsymbol{X}^i_t)_{\text{baro}}$ in [19], the activity recognition in [20].

The transition in (1) describes the system dynamics and its uncertainties as a Markov sequence, defined in terms of the transition from q_{t-1} to q_t. This is done by updating the parameters of a particle $\boldsymbol{X}^i_t$ according to possible movement patterns of a human, such as walking direction, speed, or step length, often referred to as Pedestrian Dead Reckoning (PDR). In addition, unrealistic movements like walking through walls should be restricted. For this purpose, special motion models are integrated that incorporate knowledge about the building by using a spatial representation, e.g. a graph. In this work, we use a motion model based on a triangulation mesh as presented in our previous papers [11] and [8]. The mesh models only walkable spaces, omitting any obstacle or wall. The motion model uses the IMU to calculate the current heading change and stride length, and moves the particle forward accordingly. If a particle would leave the mesh, the motion is interrupted and its weight is halved while still updating the heading. This causes particles that hit obstacles or walk the wrong way to be penalized, and after several iterations of making unrealistic movements, they are removed by resampling.

III. Particle Distribution Optimization

A. Kullback-Leibler Sampling/Resampling

Kullback-Leibler Divergence sampling (KLD), introduced by [21], aims to dynamically adjust the number of particles to improve the estimation of target distributions in uncertain scenarios (i.e. larger number of particles) while maintaining computational efficiency in more certain ones (i.e. smaller number of particles). It is grounded on the idea of calculating the Kullback-Leibler distance, in this case on could also speak of approximation error, between the current posterior $p(\boldsymbol{q}_t \mid \boldsymbol{o}_{1:t})$ and the true target $\tau(\boldsymbol{q}_{1:t})$, ensuring this error remains below a predefined threshold ϵ with a certain probability $1-\delta$. To derive ϵ it is assumed that $\tau(\boldsymbol{q}_{1:t})$ is based on a discrete, piecewise constant distribution, such as a histogram, overlayed over the multi-dimensional state space. Each bin k of the histogram is either with support, i.e., $n > 0$ particles, or without support, i.e., $n = 0$ particles. It turns out, that the required number of particles is proportional to the inverse of ϵ, and first order linear in the number k of bins with support. This means, that the true posterior is not required to be known, and it is sufficient to determine k. This can be done by estimating k by counting the number of particles that fall into an empty bin or not.

In practice, new particles are then sampled as long as the threshold ϵ is not exceeded, either while sampling [21] or resampling [22]. The adaptability of KLD lies in its response to the variance of particles; a higher variance leads to more supported bins and, consequently, a higher number of particles. In the case of CONDENSATION, i.e. the transition is chosen as importance distribution [23], using KLD while resampling has the advantage that it also considers the particle's weight, instead of only the transition function.

The implementation is straightforward and can be seen in Algorithm 1. Here, the particle filter is initialized with an a priori distribution $\mu(\boldsymbol{X}^i_1)$, often an uniform distribution of particles over the complete building, the parameters δ, ϵ, the effective sample size N_{eff} with its threshold $\hbar$, a number N of particles, the maximum number of particles N_{max} and the histogram H. From line 1 to 3 the particles filters is processed as in [8]. The resampling step begins in line 4, where in line 7 new particles are sampled based on some resampling strategy, e.g. multinomial resampling [24]. After that, the above described KLD process is carried out. N is the number of particles needed to reach our bounded ϵ and thus

Algorithm 1 CONDENSATION with KLD-Resampling

Input: A Priori $\mu(\boldsymbol{X}_1^i)$, δ, ϵ, H, N, N_{max}

for $i = 1$ **to** N **do**
  Run CONDENSATION particle filter as in [8]
end for

if $N_{\text{eff}} \leq \hbar$ **then**
  $k = 0$, $M = 0$, Empty bins b in H
  repeat
    Sample $\{\frac{1}{N}, \boldsymbol{X}_t^i\} \sim \Upsilon_t = \{W_t^i, \boldsymbol{X}_t^i\}_{i=1}^N$
    if $\boldsymbol{X}_t^i$ falls into empty bin b of H **then**
      $k = k + 1$
      $b =$ non-empty
      $N = \frac{k-1}{2\epsilon} \left\{1 - \frac{2}{9(k-1)} + \sqrt{\frac{2}{9(k-1)}}\, z_{1-\delta}\right\}^3$
    end if
    $M = M + 1$
  until $M \geq N$ or $M \geq N_{\text{max}}$
end if

updates for each particle that falls into an empty bin. While at the beginning of the processing almost all bins are empty, this leads to a continuous increase of k and therefore N. As more and more bins have support, N is updated only occasionally until M finally exceeds it, ending the update loop.

B. Kernel Density Resampling

Instead of resampling from the approximating particle set, resampling directly from the true target $\tau(\boldsymbol{q}_{1:t})$ or at least from the probability density $\tau(\boldsymbol{q}_{1:t})$ of the posterior would give a higher diversity among particles without losing focus on high probability regions. A classical method for estimating a probability density from a set of (weighted) samples is the kernel density estimator (KDE). The KDE uses a kernel function K such that $\int K(u)du = 1$, placed on each particle $\{W_t^i, \boldsymbol{X}_t^i\}$ with a given bandwidth $h_1 \dots h_d \in \mathbb{R}^+$, one for each dimension d of the state space, to approximate the target. In the context of time-sequential filtering, the kernel estimator is given by

$$\hat{\tau}(\boldsymbol{q}_t) = \frac{1}{\sum_{k=1}^N W_t^k} \sum_{i=1}^N \frac{W_t^i}{h_1 \dots h_d} \left[\prod_{j=1}^d K\left(\frac{q_{j,t} - X_{j,t}^i}{h_j}\right)\right], \tag{4}$$

where the state $\boldsymbol{q}_t$ represents the point to be estimated and the index j accesses the scalar elements of the state and particle vectors. The choice of kernel is usually not a major factor, since the kernel estimate inherits the properties of the chosen kernel. However, the bandwidth h plays a critical role in determining the quality of the estimate. For simplicity, a Gaussian kernel is often chosen [25].

The above (exact) KDE approach is very easy to implement and offers great flexibility. However, the complexity is $\mathcal{O}(NM)$ and thus several approaches to reduce the complexity of KDE have been proposed in the literature [25]–[27]. In [25] we have presented our own approach, that is so fast that it can even be used for real-time use cases, e.g. directly on a smartphone for self-localization. Based on this, we introduce kernel density resampling (KDR) to resample particles according to $\boldsymbol{X}_t^i \sim \tau(\boldsymbol{q}_{1:t})$. Unlike KLD-resampling, this approximates in a continues function, allowing new particles to be drawn at arbitrary positions within the state space. To prevent sampling in non-walkable areas not covered by the spatial model, e.g. between walls or outside the building, a validation check must be included for each newly drawn particle. If the particle has a position matching the spatial model, keep it; if not, draw another particle to keep the number of particles constant over time.

It is hoped that KDR, due to its continuous nature, will allow the reduction of sample impoverishment by resampling particles in adjacent regions, ignoring the spatial boundaries. Of course, the number of these particles depends on the chosen bandwidth h and the wall thickness. However, this may have drawbacks in situations where spatial constraints are useful, e.g. in buildings with many small rooms such as offices. Here, KDR with a large bandwidth h leads to a higher susceptibility to false measurements from absolute sensors such as Wi-Fi or BLE. In areas where there should be no particles due to pedestrian movement patterns and space limitations, they may be due to KDR and, in the worst case, have a high weight due to erroneous measurements. Basically, this means that the motion model and the resampling model are working against each other.

C. Randomized Filter Extensions

The following approaches to PDO are explicitly designed for target tracking and positioning. They provide at least some degree of general applicability, e.g., building independent, by using context and/or sensor knowledge to solve the problem of impoverishment.

The simplest approach to tackle the impoverishment problem is to randomly sample a few particles (e.g., $0.5\,\%$) in the state space while resampling. Depending on the size of the building and the number of these random drawn particles, it gives a slight chance to have particles in any area of the building. If a spatial model is used, this can be further restricted to walkable areas only. Imagine a situation where the IPS does not detect a floor change from the first floor to the second floor, all particles stay on the first floor and bypass the staircase. Since the motion model only allows for floor changes at staircases or elevators, this results in the filter losing track by being on the wrong floor with no way to correct it. To address this, randomly placing some particles on the intended floor can help the filter recover by giving these particles a high weight during state evaluation and resampling, allowing the filter to eventually converge on the correct floor.

Obviously, randomly placing particles in the state space introduces the risk of unpredictable behavior and potentially large approximation errors. In particular, if a faulty sensor, such as Wi-Fi, suggests an implausible location that would normally be ignored due to the absence of particles in that

area, the presence of a randomly drawn particle can lead to significant inaccuracies. Nevertheless, we used this approach for our IPS in the EvAAL competition [28]. It was one of the main reasons why we were able to win, since the test track included several elevator rides, which is still a hard problem to solve accurately using only the smartphone's IMU.

D. Support Filter Approach

Another approach to solving impoverishment is to combine two corresponding particle filters. A tool that allows multiple particle filters to approximate the same joint posterior to be run in parallel is the interacting multiple model particle filter (IMMPF) [29]. In [16] we suggested by defining a so-called dominant filter, which is an IPS with a restrictive motion model, configured to achieve the highest possible accuracy in position estimation, but thus highly susceptible to sample impoverishment. The second, called the support filter, uses a simple multivariate Gaussian, i.e. random direction and distance, to simulate motion and the same absolute sensor model, e.g. BLE or Wi-Fi, as the dominant one. Despite the rather uncertain/not very accurate position estimation of the support filter, it is able to recover quickly from erroneous measurements, and getting stuck is a very unrealistic scenario.

The idea of this so-called *support filter approach* (SFA) is to use the dominant filter for position estimation and the support filter to resolve possible impoverishment. This can be done by comparing their respective marginal posteriors, again using a Kullback-Leibler divergence $D_{\mathrm{KL}}(P\|Q)$ to get an idea of how much they differ. Since both posteriors are represented by a disjoint set of particles, we again use a KDE of the form (4) with a Gaussian kernel to be able to sample at arbitrary positions and thus compute $D_{\mathrm{KL}}(P\|Q)$ using the arithmetic mean between the two. If the filters are similar, the position estimation of the dominant filter should be reasonable, if not, it is very likely that one of the filters is blocked or has lost track due to impoverishment. This is solved by exchanging particles between the two filters in the mixing stage of the IMMPF. Here, a Markov transition matrix Π controls the probability that the dominant filter samples a particle from its own posterior or from the posterior of the supporting filter, and vice versa. This means, if the dominant filter loses track, diverging from the support filter, and thus $D_{\mathrm{KL}}(P\|Q)$ increases, the probability of sampling more particles from the support filter also increases.

Having Wi-Fi as absolute positioning information, a transition matrix $\Pi_{\text{KL+wifi}}$ of the IMMPF can be denoted as

$$\Pi_{\text{KL+wifi}} = \begin{pmatrix} q_{\mathrm{KL}}(D_{\mathrm{KL}}) & 1 - q_{\mathrm{KL}}(D_{\mathrm{KL}}) \\ 1 - q_{\text{wifi}}(\boldsymbol{o}_t) & q_{\text{wifi}}(\boldsymbol{o}_t) \end{pmatrix} . \tag{5}$$

where $q_{\mathrm{KL}}(D_{\mathrm{KL}}) = e^{-\lambda D_{\mathrm{KL}}}$ is an exponential function providing a probability based on $D_{\mathrm{KL}}(P\|Q)$ and q_{wifi} denotes a quality metric between 0 and 1 of the current available Wi-Fi measurements. For this, we count the number of available RSSI measurements, as for optimal multilateration there needs to be a minimum number of transmitters in the building, and we evaluate the RSSI, as lower values (e.g. $< -80\,\mathrm{dB}$) indicate that either there is no good coverage of transmitters in the vicinity, which means that the signal has to travel a longer distance and is therefore exposed to more noise factors, or there are strong attenuation effects [18]. In (5) the behavior of mutual support between the two filters is controlled. If the dominant filter suffers from sample impoverishment the probability $[\Pi_{\text{KL+wifi}}]_{12} = 1 - q_{\mathrm{KL}}(D_{\mathrm{KL}})$ increases with diverging support filters, and thus the number of particles drawn by the dominant filter from the support filter increases by $[\Pi_{\text{KL+wifi}}]_{12} \cdot 100\,\%$.

E. Uniform distributional Evaluation Resampling

A disadvantage of the previously proposed SFA is the added complexity and computing power required by the introduction of a second particle filter. For this reason, we present a novel approach called *uniform distributional evaluation* (UDE) resampling. Instead of randomly sampling some particles in the resampling step, as suggested in III-C, we instead simulate the support filter by sampling from the evaluation density $p(\boldsymbol{o}_t \mid \boldsymbol{q}_t)$. To do this, we introduce a separate static particle set $\Upsilon_{\text{static}} = \{W^j, \boldsymbol{X}^j\}_{j=1}^{M}$ that is uniformly distributed in the building. Υ_{static} is not subject to any motion model, but all M particles have a constant position that never changes, similar to an occupancy grid [7]. When resampling, the particles of Υ_{static} are weighted according to $W^j \propto p(\boldsymbol{o}_t \mid \boldsymbol{q}_t)$. We then simulate the mixing of the IMMPF by drawing a small number $\hbar_{\text{UDE}}$ of particles according to $[\frac{1}{N}, \boldsymbol{X}_t^i] \sim \Upsilon_{\text{static}}$, while the majority, $N - \hbar_{\text{UDE}}$, are sampled according to $[\frac{1}{N}, \boldsymbol{X}_t^i] \sim \Upsilon_t$, combined in the new set of particles $\dot{\Upsilon}_t$, which is used for the next filter update. Only the resampling step uses Υ_{static}, so other filter processes like state estimation do not require any modifications.

The main advantage of UDE is that there is no impoverishment, since each region of the building is uniformly populated by particles, and thus the state space is completely covered under asymptotic conditions. Furthermore, the implementation of this approach is trivial, since any existing resampling method can be extended, while the complexity of the computation hardly increases, because only a second evaluation step of the static particles has to be taken into account. Nevertheless, UDE has similar disadvantages as the random approach in III-C. While in SFA the support filter focuses on relevant areas due to the simplified motion model, UDE is susceptible to measurement errors as it considers the entire state space and may well lead to unpredictable behavior, e.g. unwanted short jumps across the building [8].

IV. Experiments

The test data was collected in two different buildings, SHL and ROT, using our open source mobile recording application [30], which measures Wi-Fi RSSI, smartphone IMU data, and barometric pressure when available. As can be seen in Fig. 1, the SHL is a university campus with approximately $12\,000\,\mathrm{m}^2$ of walkable area, consisting of two buildings connected by a large courtyard and an underground passage. The upper building is the faculty building, mainly providing office space

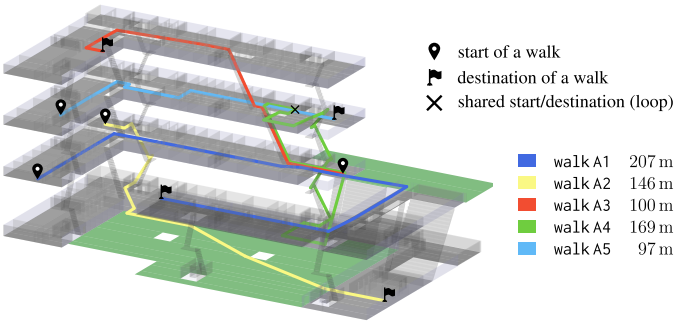

Fig. 1: The SHL campus of the Technical University of Applied Sciences Würzburg-Schweinfurt given by [31]. It has two buildings, one for faculty (top left) and one for lectures (bottom right), and a total accessible area of about $12\,000\,\mathrm{m}^2$. The facade is made of metallized glass panels and the interior is a mixture of reinforced concrete and drywall.

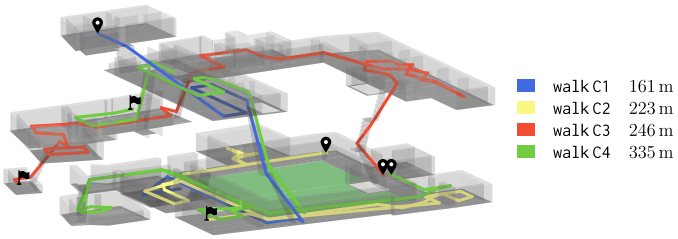

Fig. 2: ROT was built in the 13th century and has $2000\,\mathrm{m}^2$ walkable space. It is made of thick stone walls on the ground floor and wood framing on the upper floors. Figure taken from [31].

for the university's staff, having four floors with a footprint of $77\,\mathrm{m} \times 50\,\mathrm{m}$. The lower building consists of lecture halls and a cafeteria with an area of $110\,\mathrm{m} \times 60\,\mathrm{m}$. We conducted several recordings of 5 walks throughout the building using 9 different commercial available smartphones with a subject group of 4 people. There are 34 access points distributed throughout the campus, using both $2.4\,\mathrm{GHz}$ and $5\,\mathrm{GHz}$. The second building ROT is illustrated in Fig. 2, having 4 walks recorded using 3 different smartphones with a subject group of 3 people. It was built in the 13th century, undergone several reconstructions and nowadays serves as a museum. It measures $70\,\mathrm{m} \times 50\,\mathrm{m}$, has $2000\,\mathrm{m}^2$ of walkable space and is made of thick stone walls on the ground floor and wood framing on the upper floors. As the building has no Wi-Fi infrastructure, we installed 42 WEMOS D1 mini (ESP8266) microcontroller as Wi-Fi beacons on available electrical outlets.

To capture the ground truth, we place numbered markers with a unique color for each walk inside the building. The positions of these markers are measured and added to the spatial model. While recording the test walks, the user simply clicks a button on the smartphone application when passing a ground truth marker. This saves a timestamp to the recording file, which can then be used to assume a constant speed of movement between two successive markers, allowing a position and timestamp to be interpolated between them. This is obviously not a perfect way to measure, but if you use enough markers to avoid unrealistic interpolations, e.g. through walls, it can be fair enough for error measurements. To finally compute the approximation error, we compare the interpolated ground truth position with the current estimate. Similar to [9], errors in the z-direction are penalized by tripling the z-value.

SHL					ROT			
1	2	3	4	5	1	2	3	4
8 %	44 %	0 %	8 %	11 %	20 %	6 %	8 %	7 %

TABLE I: Percentage of 50 simulation runs that got stuck and did not reach the destination within a radius of $10\,\mathrm{m}$ for all conducted walks of the data sets. It uses the particle filter as described in II without any PDO approaches.

The main purpose of this work is to investigate the performance of PDO approaches when faced with the problem of impoverishment. Therefore, we first calculated the percentage out of 50 simulation runs for each walk of all datasets that got stuck and did not reach the destination within a radius of at least $10\,\mathrm{m}$. The results are shown in TABLE I. They are computed using the particle filter described in II with 5000 particles, activity detection, barometer and Wi-Fi RSSI. In the following, we will refer to it as None because it does not use any PDO approaches. In summary, the probability of missing is $13.86\,\%$ for the SHL campus and $10.25\,\%$ for ROT. Both of these values could be detrimental in real-world scenarios and reduce user acceptance.

In the following we will concentrate on the two worst results, walk 2 in SHL and walk 1 in ROT. Walk 2 starts only a few meters away from a stairwell that is heavily shielded by reinforced concrete and where only highly attenuated RSSI values can be measured. When entering the stairwell, the uncertainty is still very high due to the fact that the particles are uniformly distributed throughout the building instead of having a known initial position. This often results in the upcoming floor change not being accurately detected, causing the system to get stuck as jumping to another floor is limited by the motion model. Another challenging area is the courtyard between the two buildings, where there are no outdoor Wi-Fi access points, so the system can only receive signals from inside the building. This is not optimal as the signal passes through metalized windows, which are known for multipath effects and attenuation [18]. Walk 1 in ROT has a $20\,\%$ chance of getting completely stuck due to erroneous Wi-Fi measurements between $10\,\mathrm{s}$ and $20\,\mathrm{s}$. Even for paths that reach the destination, this leads to a very high error until the system finally converges after $60\,\mathrm{s}$ (see Fig. 4, red line). For a user, this may be too late, as the behavior is not obvious.

The overall results for all presented PDO approaches are shown in TABLE II. We will now discuss them in the order in which they were introduced.

As the KLD approach uses an adaptive size of particles, we tested to variations, KLD 5K having $N = 5000$ particles with a maximum of $N_{\mathrm{max}} = 10000$ particles that can be

	SHL walk 2				ROT walk 1			
	$\bar{x}$	$\tilde{\sigma}$	$\tilde{x}_{75}$	stuck	$\bar{x}$	$\tilde{\sigma}$	$\tilde{x}_{75}$	stuck
None	19.09	19.23	31.77	44 %	9.30	9.76	13.41	20 %
KLD 5k	18.88	16.18	29.43	40 %	10.12	9.57	15.39	18 %
KLD 25k	16.86	21.57	26.44	32 %	6.26	7.38	5.72	8 %
KDR	27.63	20.82	47.26	100 %	14.30	10.68	22.70	40 %
Random	18.81	19.07	30.89	20 %	2.88	2.34	3.40	0 %
SFA	5.27	3.65	7.25	2 %	2.76	2.53	3.31	0 %
UDE	13.94	15.87	22.29	0 %	2.80	2.30	3.31	0 %

TABLE II: Overall positioning results in m for walk 2 of SHL and walk 1 of ROT using the here presented PDO approaches. We ran 50 simulation runs with 5000 particles. A run got stuck if it did not reach the target within a radius of at least 10 m.

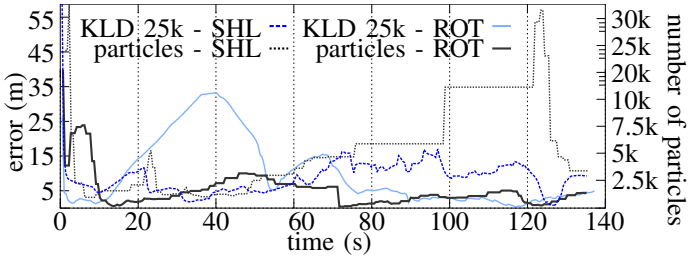

Fig. 3: Error (blue) and number of particles (black) over time using KLD for walk 2 in SHL (dotted) and walk 1 in ROT (solid). Only a very weak correlation between both entities can be observed. For SHL the number of particles increases significantly when going into the large outside open space (courtyard) between 80 s and 120 s.

reached and KLD 25k with $N = 25000$ and $N_{\text{max}} = 50000$. As suggest by [21] we set the error bound to $\delta = 0.01$ and $\epsilon = 0.05$. Looking again at TABLE II, only KLD 25k is able to improve the results of the standard IPS None due to its larger number of particles. It is important to note that there is no linear relationship between increasing particle number and decreasing error [17], thus simply maximizing the number of particles does not automatically lead to convergence. The relationship between errors and the adaptive behavior of KLD is illustrated in Fig. 3. Initially, a weak correlation is observed, as high errors from faulty Wi-Fi in ROT between 20 s and 40 s don't significantly affect particle numbers, unlike in SHL's courtyard between 80 s and 120 s, where particle numbers increase despite smaller errors. This discrepancy is related to the KLD approach's assumption of a discrete piecewise constant distribution, e.g. a grid, for the true posterior. The required number of particles is linearly related to the number of bins with support. In SHL's open courtyard, particle scattering increases due to the lack of obstacles, leading to a higher particle number. In contrast, the confined spaces of ROT with obstacles lead to a slower particle distribution and a lower number of particles. Reducing the histogram bin size increases the number of particles, e.g. 10 cm bins require on average 45554 particles, but questions the efficiency of the KLD approach as it aims to minimize particle usage while maintaining accuracy. Adapting the motion model or the KLD approach could improve results in complex environments.

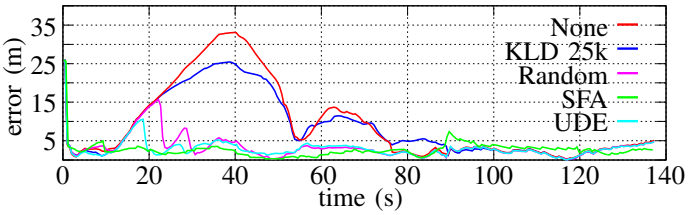

Fig. 4: Error over time using the PDO approaches improving None for a single recording of walk 1 in ROT using the same random seed. Between 10 s and 20 s the Wi-Fi signal was highly attenuated, causing an impoverishment and thus a high error.

KDR underperforms in the given scenarios due to the limitations imposed by the spatial model, which restricts particle movement to walkable regions, essentially allowing movement only on floors or stairs, and excluding any three-dimensional positioning such as being "in the air". This results in a bottleneck at staircases, where particles become densely packed due to the restricted vertical movement, while many particles remain on the floor. KDR's KDE approximation over this particle set poorly defines the probability density along the vertical axis due to the constraints of the spatial model, resulting in resampling that fails to draw particles on staircases, effectively preventing floor changes in all cases. For two-dimensional scenarios, KDR successfully eliminates particle impoverishment, reducing the rate of stuck runs to 0 %. However, it does not significantly improve the positioning error for successful runs. Despite its current limitations, the approach shows potential, especially with its analytical representation of the posterior. Resampling strategies that focus more on ground changes, such as adapting KDR for two and a half dimensional spatial models, may offer improvements.

In TABLE II, Random, SFA, and UDE show significant improvements, especially in ROT where their performance is nearly identical. In Fig. 4 the error over time is plotted for a single recording of the data set, using a consistent random seed to ensure comparability. The high error due to impoverishment can be clearly seen in the case of None (red), with all methods converging to an error below 5 m after 80 s at the latest. Except for SFA (green), which uses IMMPF and produces a distinct error curve, the other methods show similar error patterns. Random (pink) and UDE (teal) suffer briefly from the erroneous Wi-Fi measurements between 10 s and 40 s, but recover, with UDE improving faster because it is less dependent on the random factor. In larger buildings, the effectiveness of random resampling diminishes, requiring more draws to counteract inertia but risking uncontrolled behavior due to measurement errors. This is evident in walk 2 in SHL, shown in Fig. 5, where Random (pink) struggles with Wi-Fi quality issues in the stairwell and courtyard, leading to inaccurate positioning after 60 s. Only after entering the lower building the weighting of the random particles is sufficient to recover from impoverishment. UDE (teal), with its static particles in the courtyard, maintains its trajectory

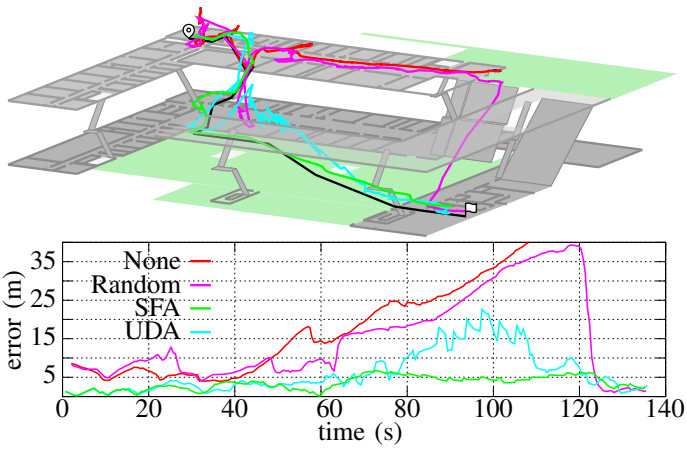

Fig. 5: Trajectory (top) and error over time (bottom) of walk 2 of SHL using the same random seed. KDE and KLD also get stuck in the first stairwell, similar to None, and are therefore omitted. In the large open space (courtyard) between 80 s and 120 s the error increases significantly because Wi-Fi access points are only available inside the buildings, so the signals are highly attenuated.

outdoors despite high error rates, showing some resilience. They significantly outperform the None approach in reducing impoverishment effects, with UDE eliminating them entirely in evaluated walks. However, they are highly dependent on the quality of the sensor model, such as the Wi-Fi RSSI, for accurate positioning. Both methods provide viable solutions for limiting impoverishment in indoor scenarios, with Random being particularly resource efficient.

Finally, the SFA method, utilizing multi-modal fusion with IMMPF, shows excellent performance across all experiments. In ROT (see Fig. 4), SFA (green) effectively mitigates initial high errors by relying less on the support filter and thus the Wi-Fi sensor model during periods of low Wi-Fi quality (between 10 s and 20 s). Instead, the restrictive motion model with its PDR of the dominant filter is prioritized for position estimation. Nevertheless, there is a notable increase in error during floor changes at 90 s and just before 120 s, where the open staircase layout leads to higher Wi-Fi quality because a larger number of RSSI measurements are received from unique transmitters of the two adjacent floors. This results in increased uncertainty because no transmitters are placed directly on the stairs, and a signal received from further away typically has more noise.

In SHL (see Fig. 5), SFA demonstrates superior performance in the courtyard around 70 s despite an increase in average positioning error. This increase is attributed to both the poor Wi-Fi quality in the open area and strong drift, exacerbated by obstacles that induce bimodality. However, SFA maintains a lower error rate than UDE (teal) due to a slightly higher influence of the PDR over the uncertain Wi-Fi, leading to better overall performance in this scenario. The observations suggest that while SFA effectively reduces error in challenging environments by prioritizing reliable data sources, its performance is contingent on the quality and availability of sensor data. In extended outdoor scenarios, results between SFA and UDE could converge, highlighting the need for integrating additional sensor inputs, such as GNSS or outdoor Wi-Fi, for improved positioning in real-world applications.

V. Conclusion

This work focused on an experimental comparison of different PDO approaches to improve smartphone-based IPS. In particular, we aimed to mitigate sample impoverishment - a persistent challenge in particle filtering. The PDO approaches were tested under realistic conditions in two different environments to evaluate their impact on accuracy and reliability.

For the data-driven approaches, namely Random, SFA, and UDE, the experimental results showed a significant improvement in positioning accuracy and system robustness in both environments. In particular, SFA provided a valuable mechanism for system recovery in scenarios where either restrictive motion models or erroneous sensor models would lead to impoverishment. Despite being inferior to SFA in terms of position estimation accuracy, UDE is simple, fast to implement, and has low computational complexity. In contrast, KLD and KDR, which promise general applicability, i.e. independent of the sensor models used in the IPS, could not significantly improve the overall results in practice, despite their good theoretical prospects.

In future work, we plan to further investigate the capabilities of SFA by searching for general quality metrics that can make reliable statements independent of the sensors used. We also believe that the KDR approach has great potential due to its analytical properties, e.g. finding the true mode of a multimodal distribution. In order to do this, the KDR must be adapted to the restrictive motion models in use.

References

[1] Q. Wang, M. Fu, J. Wang, *et al.*, "Recent Advances in Pedestrian Inertial Navigation Based on Smartphone: A Review," *IEEE Sensors Journal*, vol. 22, no. 23, pp. 22 319–22 343, Dec. 2022.

[2] Y. Gu, A. Lo, and I. Niemegeers, "A Survey of Indoor Positioning Systems for Wireless Personal Networks," *IEEE Communications Surveys and Tutorials*, vol. 11, no. 1, pp. 13–32, 2009.

[3] P. Davidson and R. Piché, "A Survey of Selected Indoor Positioning Methods for Smartphones," *IEEE Communications Surveys and Tutorials*, vol. 19, no. 2, pp. 1347–1370, Apr. 2017, Publisher: Institute of Electrical and Electronics Engineers Inc.

[4] T. Savić, X. Vilajosana, and T. Watteyne, "Constrained Localization: A Survey," *IEEE Access*, vol. 10, pp. 49 297–49 321, 2022.

[5] F. Gu, X. Hu, M. Ramezani, *et al.*, "Indoor Localization Improved by Spatial Context;A Survey," *ACM Computing Surveys*, vol. 52, no. 3, 64:1–64:35, Jul. 2019.

[6] X. Guo, N. Ansari, F. Hu, Y. Shao, N. R. Elikplim, and L. Li, "A Survey on Fusion-Based Indoor Positioning," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 1, pp. 566–594, 2020.

[7] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA, USA: The MIT Press, 2005.

[8] T. Fetzer, M. Bullmann, M. Ebner, S. Kastner, F. Deinzer, and M. Grzegorzek, "Interacting Multiple Model Particle Filter for Indoor Positioning Applications," en, in *Proceedings of the 2023 International Technical Meeting of The Institute of Navigation*, Long Beach, California, 2023, pp. 1089 –1100.

[9] *Evaluating aal systems through competitive benchmarking*, https://evaal.aaloa.org/, Accessed: 2023-03-12.

[10] F. Ebner, T. Fetzer, L. Köping, M. Grzegorzek, and F. Deinzer, "Multi Sensor 3D Indoor Localisation," in *Indoor Positioning and Indoor Navigation (IPIN), Int. Conf. on*, Banff, Canada: IEEE, 2015.

[11] T. Fetzer, F. Ebner, M. Bullmann, F. Deinzer, and M. Grzegorzek, "Smartphone-Based Indoor Localization within a 13th Century Historic Building," *Sensors (Basel, Switzerland)*, vol. 18, no. 12, pp. 1–32, 2018.

[12] M. Ebner, T. Fetzer, M. Bullmann, S. Kastner, F. Deinzer, and M. Grzegorzek, "PIPF: Proposal-Interpolating Particle Filter," in *2022 IEEE 12th International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, ISSN: 2471-917X, Sep. 2022, pp. 1–8.

[13] S. Knauth, "Smartphone PDR positioning in large environments employing WiFi, particle filter, and backward optimization," in *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, ISSN: 2471-917X, Sep. 2017, pp. 1–6.

[14] A. Moschevikin, A. Galov, A. Soloviev, A. Mikov, A. Volkov, and S. Reginya, "Realtrac Technology Overview," in *International Competition on Evaluating AAL Systems through Competitive Benchmarking*, Springer, 2013, pp. 60–71.

[15] S. Park, J. H. Lee, and C. G. Park, "Robust Pedestrian Dead Reckoning for Multiple Poses in Smartphones," *IEEE Access*, vol. 9, pp. 54 498–54 508, 2021.

[16] T. Fetzer, F. Ebner, F. Deinzer, and M. Grzegorzek, "Recovering from sample impoverishment in context of indoor localisation," in *2017 International Conference on Indoor Positioning and Indoor Navigation*, 2017, pp. 1–8.

[17] Z. Chen, "Bayesian Filtering: From Kalman Filters to Particle Filters, and Beyond," *Statistics*, vol. 182, no. 1, pp. 1–69, 2003.

[18] F. Ebner, T. Fetzer, F. Deinzer, and M. Grzegorzek, "On Wi-Fi Model Optimizations for Smartphone-Based Indoor Localization," *ISPRS International Journal of Geo-Information*, vol. 6, no. 8, 2017.

[19] T. Fetzer, F. Ebner, F. Deinzer, and M. Grzegorzek, "Using Barometer for Floor Assignation within Statistical Indoor Localization," en, *Sensors*, vol. 23, no. 1, p. 80, Jan. 2023, Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.

[20] M. Ebner, T. Fetzer, M. Bullmann, F. Deinzer, and M. Grzegorzek, "Recognition of typical locomotion activities based on the sensor data of a smartphone in pocket or hand," *Sensors*, vol. 20, no. 22, 2020.

[21] D. Fox, "KLD-Sampling: Adaptive Particle Filters," in *Advances in Neural Information Processing Systems*, vol. 14, MIT Press, 2001.

[22] T. Li, S. Sun, and T. P. Sattar, "Adapting sample size in particle filters through KLD-resampling," en, *Electronics Letters*, vol. 49, no. 12, pp. 740–742, Jun. 2013, Publisher: IET Digital Library.

[23] M. Isard and A. Blake, "CONDENSATION - Conditional Density Propagation for Visual Tracking," *Int. Journal of Computer Vision*, vol. 29, no. 1, pp. 5–28, 1998.

[24] R. Douc and O. Cappé, "Comparison of resampling schemes for particle filtering," in *Image and Signal Processing and Analysis, 2005. ISPA 2005. Proc. of the 4th Int. Symp. on*, IEEE, 2005, pp. 64–69.

[25] M. Bullmann, T. Fetzer, F. Ebner, F. Deinzer, and M. Grzegorzek, "Fast Kernel Density Estimation Using Gaussian Filter Approximation," in *International Conference on Information Fusion, FUSION 2018*, Cambridge, United Kingdom, 2018, pp. 1233–1240.

[26] T. A. O'Brien, K. Kashinath, N. R. Cavanaugh, W. D. Collins, and J. P. O'Brien, "A fast and objective multidimensional kernel density estimation method: fastKDE," *Computational Statistics and Data Analysis*, vol. 101, pp. 148–160, Sep. 2016.

[27] M. P. Wand, "Fast Computation of Multivariate Kernel Estimators," *Journal of Computational and Graphical Statistics*, vol. 3, no. 4, pp. 433–445, Dec. 1994.

[28] J. Torres-Sospedra, A. R. Jiménez, S. Knauth, *et al.*, "The smartphone-based offline indoor location competition at IPIN 2016: Analysis and future work," *Sensors (Switzerland)*, vol. 17, no. 3, 2017.

[29] Y Boers and J. N. Driessen, "Interacting multiple model particle filter," *October*, vol. 150, no. 5, pp. 344–349, 2003.

[30] M. Bullmann, F. Ebner, M. Ebner, T. Fetzer, and L. Köping, *SensorReadoutApp*, https://github.com/simpleLoc/SensorReadoutApp, Accessed: 14.03.2024), 2020.

[31] F. Ebner, *Smartphone-Based 3D Indoor Localization and Navigation*, en. Logos Verlag Berlin GmbH, Jan. 2021, Google-Books-ID: iKsSEAAAQBAJ.