ARTICLE

Exploring the Limits of PDR-based Indoor Localisation Systems under Realistic Conditions

Robert Jackermeier and Bernd Ludwig

Chair for Information Science, University of Regensburg, Regensburg, Germany

ABSTRACT

Pedestrian Dead Reckoning (PDR) plays an important role in many (hybrid) indoor positioning systems since it enables frequent, granular position updates. However, the accumulation of errors creates a need for external error correction. In this work, we explore the limits of PDR under realistic conditions using our graph-based system as an example. For this purpose, we collect sensor data while the user performs an actual navigation task using a navigation application on a smartphone. To assess the localisation performance, we introduce a task-oriented metric based on the idea of landmark navigation: instead of specifying the error metrically, we measure the ability to determine the correct segment of an indoor route, which in turn enables the navigation system to give correct instructions. We conduct offline simulations with the collected data in order to identify situations where position tracking fails and explore different options how to mitigate the issues, e.g. through detection of special features along the user's path or through additional sensors. Our results show that the magnetic compass is often unreliable under realistic conditions and that resetting the position at strategically chosen decision points significantly improves positioning accuracy.

KEYWORDS
Pedestrian Navigation, Indoor User Localisation, Inertial Navigation,
Map-Matching

1. Introduction

Indoor positioning for pedestrian navigation is still an open problem as no positioning system that delivers absolute – such as GPS – coordinates is available. Many solutions providing precise (sub meter) localisation require additional technical devices (Guo et al. 2015; Pham and Suh 2016; Romanovas et al. 2013). However, they are not at disposal in everyday life situations at which pedestrian navigation systems target. Technically simpler solutions use sensors that come with every smartphone, such as accelerometers, gyroscopes, and step counters (Basso, Frigo, and Giorgi 2015; Verma et al. 2016). Such approaches provide relative positioning data, and often suffer from cold start problems (Harle 2013). In the remainder of this section, we discuss systematic issues that limit the capability of smartphone sensors to reliably estimate indoor positions even after a series of sensor updates. Furthermore, we identify the fact that
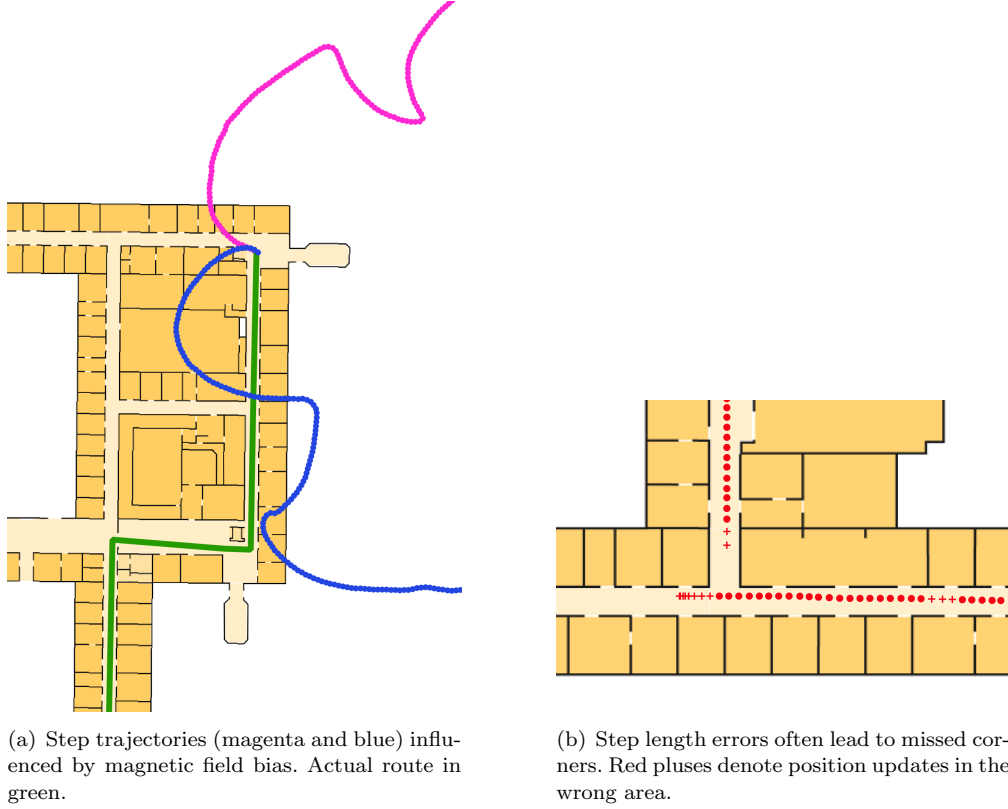
CONTACT Robert Jackermeier. Email: robert.jackermeier@ur.de

(a) Step trajectories (magenta and blue) influenced by magnetic field bias. Actual route in green.

(b) Step length errors often lead to missed corners. Red pluses denote position updates in the wrong area.

Figure 1. Different types of bias that limit indoor positioning.

users – as opposed to technical devices that navigate autonomously in physical spaces – have to cooperate with a navigation system and thereby perform activities that cannot be observed directly as a source for noisy indoor positioning.

1.1.   *Limits of PDR-based Indoor Localisation*

As will be detailed in Sec. 2, it is common to apply inertial sensors to localise users in smartphone-based navigation systems for everyday usage. Unfortunately, magnetic and electrostatic fields cause noisy data and thereby limit the quality of measurements. The graphic in Fig. 1(a) illustrates the limit caused by *magnetic bias*. The graphic shows two extreme logs: although in both cases the test person walked on the green line, the smartphone's orientation sensor computes wrong data from which the actual orientation of the test person cannot be reconstructed. Furthermore, the geometry of the green line is hardly recognisable.

The graphic in Fig. 1(b) illustrates the limit of *step length*. The test person walks along the corridor – each step is represented by a red dot. At the corner, the person turns to the right. However, as the smartphone estimated the step length incorrectly, the indoor positioning system assumes the test person to still walk straight on although the turn has already been completed (red pluses).

Solutions based on WiFi receivers can compute absolute coordinates of an area in which a user is located. Therefore, they are useful for observing positions, but fail in tracking a user's movement precisely (Waqar, Chen, and Vardy 2016). In fact, with a standard smartphone WiFi signal updates are slow and often noisy due to unknown

conditions of the physical environment. As a consequence, small movements of a walking person are hard to detect reliably. Detecting large movements only leads to a huge lag in tracking movements as illustrated in Fig. 3(b) (see the yellow line). Further typical types of movements a walking person performs regularly such as turning around a corner cannot be observed directly as the physical principle of estimating positions using radio signals does not allow for it.

Combining sensors for movement and position could – in conclusion – improve the quality of indoor positioning. In fact, as the green line in Fig. 3(b) indicates, this is true, but only to a certain degree. The limit of *delay* still prevents higher accuracy. In buildings with many corners, this limit may pose an additional problem to any positioning algorithm: In the context of pedestrian navigation we are investigating in this paper, unreliable position data leads to ambiguous interpretations. Navigation systems assist users to walk along a route calculated earlier and runs into trouble if position data is uncertain resulting in high probabilities that the system assumes users to take paths they actually have not taken.

From such misinterpretations, severe confusions may arise that are well-known even in outdoor areas. Using a navigation system as a pedestrian in a historic city centre is often frustrating as the system receives wrong or even no GPS updates and in consequence locates users at wrong positions and computes erroneous navigation instructions that users are unable to understand.

Given this situation, we are convinced that – even if in the last years many researchers worked on this topic – it is still worth new research efforts. Actually, they seem even to be mandatory for building indoor navigation systems that can provide reliable assistance to their users.

1.2. *Tracking Activities of Pedestrians*

Experience gained in the last years shows that – due to the limits described above – increasing the accuracy of state-of-the-art indoor positioning algorithms for smartphone users (which we focus on in this paper) with more data only will hardly result in sufficiently small error rates. Although – as Waqar, Chen, and Vardy (2016) point out – the mentioned state-of-the-art approaches can advantageously be used to implement indoor positioning rapidly at any location that provides the necessary infrastructure, more context information is necessary to overcome the limits described above.

For navigation systems, such context information can be derived from the calculated path to a target and the activities that persons have to perform in order to reach the target. In fact, it is just these activities that have to be observed – either directly or indirectly by calculating a sequence of position estimates.

We address exactly this issue and describe an approach that models spatial knowledge for routes with an indoor navigation graph (see Fig. 2 as an example). The graph contains all routing decisions and path segments between any two arbitrary locations in an indoor/outdoor environment. Our algorithm for tracking the activities of pedestrians snaps position updates from smartphone sensors to particular edges of the graph while they are navigated along a route calculated in advance. Position updates are computed using PDR based on a Step and Heading System (SHS) approach. Typical activities of pedestrians are:

- walking straight ahead
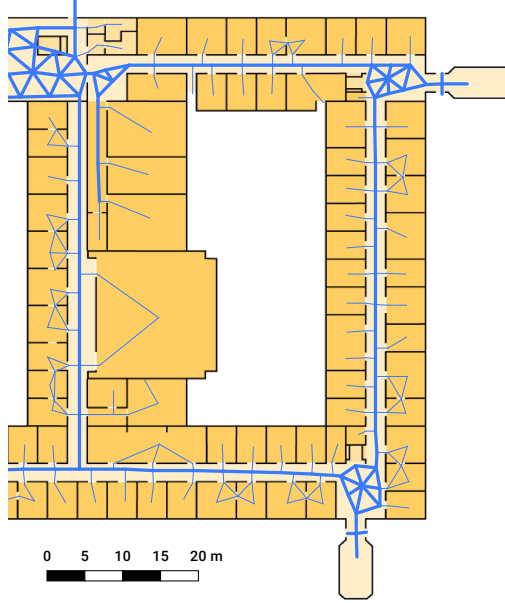- turning left/right
- walking upstairs/downstairs

Figure 2. Indoor navigation graph in one of the test areas. Main edges used for localisation are shown as thick lines. Notice the mesh-like topology in larger open areas.

- going up/down in an elevator
- taking an escalator uphill/downhill

In this work, we explore to what degree the activities of *walking straight ahead*, *turning left*, and *turning right* can be observed with a standard SHS approach. We identify situations of different accuracy rates and argue that each activity needs a particular classifier in order to increase the overall accuracy of indoor positioning.

1.3. *Tracking Pedestrians on a Navigation Route*

In order to observe activities with a standard SHS approach we map each activity on the area that has to be traversed in order to complete it. In other words, to each activity during a navigation task, there exists

(1) a corresponding navigation instruction and
(2) an area in which the activity takes place.

Therefore, we segment a complete route according to the navigation instructions calculated by the route planner of our system into a single area for each instruction and observe the corresponding activity by mapping position updates onto the area. In order to be able to evaluate the success of the approach, we introduce the **area match score** as a new metric for indoor positioning: we calculate the probability of a position update to fall into the area corresponding to the area for the current navigation instruction. In this way, we establish a task-based concept of context for indoor positioning that has the potential to overcome the limits described above.

For the evaluation, we have collected a large corpus of indoor data under realistic conditions: test persons had to walk on routes of roughly 800 m containing all types of activities several times and show quite different architectural characteristics. Therefore, the routes are typical for real world applications and avoid memory biases that could influence the walking (and information) behaviour of the test persons.

From results of the analysis, we conclude that effective positioning algorithms for pedestrian navigation must be able to apply different techniques for sensing the user's environment in order to correctly track the user activities during a navigation process.

1.4. *Structure of the Paper*

In this paper, we first report the relevant state of the art, then we explain how we relate landmark-based navigation and indoor positioning and develop our mathematical model for the posed problem. Next, we present three empirical evaluations for navigation tasks that required test persons to walk on indoor routes of varying complexity across several buildings on the campus of our university. Finally, we discuss the obtained results in the light of our task-based performance metric. We analyse limitations of the approach and derive relevant issues of future work from the insights obtained from the evaluation.

2.   State of the Art in Indoor Positioning

As already mentioned, many localisation techniques based on different types of sensors have been proposed for pedestrian indoor navigation systems and indoor positioning in a broader sense. Despite all these research efforts, there is still no technology established as a widely accepted state of the art similarly to GPS for outdoor areas. In particular, published research results – the most recent ones (e.g. from the last IPIN conference) will be discussed below – are based on data collected in short-distance and short-time experiments and therefore do not reflect typical characteristics of realistic indoor applications such as trying to reach a target in large buildings such as airports or train stations, visiting a museum, or searching an office in a multi-level building. Therefore, in order to build up a corpus containing more complex data we conducted a series of experiments that will be described in detail in a later section.

As our work is focused on positioning algorithms that assist users in everyday navigation tasks and may not need sensors beyond those available in a standard smartphone, the following review of the state of the art on PDR leaves aside more exotic approaches that require special sensors or hardware.

WiFi-based indoor localisation can be widely deployed in modern buildings where a sufficient WiFi infrastructure is usually available, but suffers from multiple problems, as Davidson and Piché (2016) point out: Creation and maintenance of radio maps is time consuming and therefore expensive. A low scan rate on current smartphones leads to disjointed position estimates. Furthermore, device heterogeneity, influence of the smartphone's orientation, and the attenuation of signals by humans are identified as disadvantages. Due to these issues, WiFi-based systems generally achieve an accuracy of at most a few meters and are suited to determine the approximate position, but not for continuous tracking.

Our own findings confirm these claims: Figure 3(a) shows some of the results from an earlier study, where the location reported by Fraunhofer's WiFi-based awiloc system[1] is wandering around in an indoor area even if the test person is not moving. Consequently, the RMSE of the position reaches up to 5 meters. Even more problems arise when the test person is moving (an example can be seen in Fig. 3(b)), where the location updates usually are lagging behind and do not match the path that was actually taken.
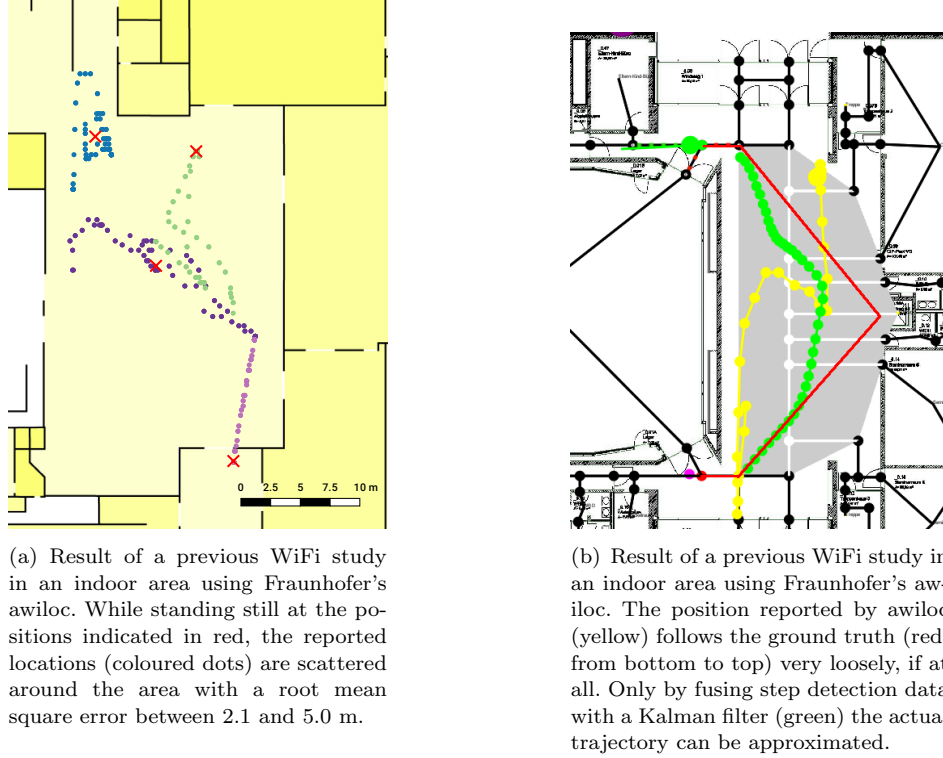
---

[1]https://www.iis.fraunhofer.de/en/ff/lv/lok/tech/feldstaerke/rssi/tl.html

(a) Result of a previous WiFi study in an indoor area using Fraunhofer's awiloc. While standing still at the positions indicated in red, the reported locations (coloured dots) are scattered around the area with a root mean square error between 2.1 and 5.0 m.

(b) Result of a previous WiFi study in an indoor area using Fraunhofer's awiloc. The position reported by awiloc (yellow) follows the ground truth (red, from bottom to top) very loosely, if at all. Only by fusing step detection data with a Kalman filter (green) the actual trajectory can be approximated.

Figure 3. Performance of WiFi signals in indoor positioning tasks.

On the other hand, Bluetooth Low Energy (BLE) beacons as another wireless localisation technique are designed to be more accurate, but are far less widespread and therefore more expensive to deploy. In particular, as they are mounted at fixed positions and send signals over a small distance only, persons could be tracked continuously only if beacons were mounted along all paths persons can walk on.

A recent development for getting an rough estimate of the user's position bases on sending a sound signal via the smartphone's loudspeaker and recording it immediately with the microphone. Rooms have particular acoustic characteristics that can be recognised to identify in which room out of a set of trained rooms the smartphone is currently located (see Rossi et al. 2013).

Both BLE and acoustic approaches are examples for zone-based localisation approaches: their aim is to identify positions in a rough approximation of an existing physical space instead of tracking user activities as we intend to do. To the best of our knowledge, there is no evidence that with these approaches user activities taking place in a very limited space may be observed reliably. For example, while it is possible to detect that a user is positioned in a foyer, it is hard to observe that he is turning right in this foyer in order to reach the door of an elevator. Consequently, zone accuracy as presented e.g. in Pulkkinen and Verwijnen (2015) has to be distinguished from our area match score: while for zone accuracy it is sufficient to locate the user in a certain zone, for the area match score, the same observation is necessary, but not sufficient as e.g. a right turn in the foyer, and in particular in front of the elevator door has to be detected if this activity is currently expected by the navigation system to continue on a computed route. Additionally, we want to note that the turn to the right may be detected by other approaches than computing a sequence of several small position changes and then reconstructing the turn from some kind of interpolation of the posi-

tion updates. For example, the accelerometer of the smartphone could help to detect a rotation that is expected for a turn.

In summary, several approaches exist that provide a rough estimate of the user's current position, but not of the user's movement. A notable exception for moving up/down is the SemanticSLAM algorithm described in Abdelnasser et al. (2016). However, users also perform other activities. For their recognition in a hybrid approach, correspondences between activities, positions, and areas can be exploited to collect redundant data reliably. Later in this paper, we will provide empirical evidence that redundant data is actually necessary as data of a single sensor often leads to ambiguous interpretations (see the discussion in Sec. 1.1).

Still the most important approach for PDR is a variant of so-called Step and Heading Systems, that detect the user's steps and try to estimate their length and direction (Harle 2013). Step detection on smartphones is historically achieved through the accelerometer using various techniques (see Susi, Renaudin, and Lachapelle 2013; Muro-de-la Herran, Garcia-Zapirain, and Mendez-Zorrilla 2014; Sprager and Juric 2015). Lately, dedicated step detector sensors are available in more and more devices. The heading can be inferred from the magnetic compass and gyroscope of smartphones, while step lengths can be either assumed as fixed or dynamic, e.g. based on the frequency (Harle 2013).

A further limit of PDR is the need for an *initial position* from which the relative positioning can start as SHS by their nature cannot compute absolute positions. Furthermore, the *positioning error increases* over time due to noisy sensor data. Given both of these problems either error correction through external sensors or an algorithm that matches the sensor data to a final position estimate are necessary to employ dead reckoning for more complex tasks such a navigating a user or other location-based services.

As in our work external sensors should be avoided, matching algorithms that provide context information are the only option to address the PDR limits. Context information is often provided in terms of maps, which in turn are often represented as discrete graphs (see e.g. Thrun, Burgard, and Fox 2005) and have been used successfully to locate robots in complex environments. For pedestrian indoor localisation, graph models of the environment were first introduced by Liao et al. (2003) in combination with the particle filtering method in order to make position estimation more robust and efficient. Since then, other researchers have adapted and improved this approach (e.g. by adding multiple sensor modalities): the system recently presented by Hilsenbeck et al. (2014) is operating on a graph generated from a 3D model of the environment. Herrera et al. (2014) use existing material from OpenStreetMap and enrich it with information about the indoor areas of a building. A similar approach is taken by Link et al. (2013), who use sequence alignment algorithms to match detected steps with the expected route. Ebner et al. (2015) generate a densely connected graph from the floor plan of a building. All these approaches have in common that creating a map is either time consuming or expensive (due to the need for special hardware) or relies on existing data. Furthermore, normally the resulting graphs do not contain any information besides the geometry of the building, making them unsuitable to use as data source for the path planner of a navigation system.

In our approach instead, we use graphs that represent all activities a user can perform in a given physical environment and contain information where these activities can be executed. The main advantage is – as already outlined in the discussion of BLE above – that we now can observe activities indirectly by reconstructing them from position updates – thereby reducing the impact of noise –, or directly by other sensing

strategies, or using hybrid approaches that fuse data from several sources.

Particle filtering has become the de facto standard for hybrid localisation systems that combine PDR and additional sensors, with many improvements proposed since its introduction. Most prominent among them is the Backtracking Particle Filter (Klepal, Beauregard et al. 2008; Beauregard, Klepal et al. 2008), which is particularly suited to generate smooth and coherent trajectories when calculations can be performed offline, but provides little or even no advantage in real-time scenarios. It does however serve as an example for map-matching that is not based on a graph structure, but on the actual floor plans. Compared to the graph-based systems described above, the environment can be represented more faithfully, with the downside of higher computational cost.

As far as the evaluation of indoor positioning is concerned, the state of the art can be surveyed best by looking at recent competitions that aim to compare the performance of indoor positioning systems. Held regularly, they provide an opportunity to gain insights into established evaluation methods. Potortì et al. (2015) report the results of the EvAAL-ETRI competition held in conjunction with the IPIN 2015 conference. To assess the error of the participating systems, they add a penalty for wrongly detected floors or buildings to the actual positioning error. The final ranking is determined by the 75% quantile of the resulting errors. In their evaluation of the 2015 EvAAL-ETRI WiFi fingerprinting competition, Torres-Sospedra et al. (2017a) rank the competitors by the mean positioning error, which is also used by Lymberopoulos et al. (2015) for the participants of the 2014 Microsoft Indoor Localization Challenge. Interestingly however, they remark that the mean error or other commonly used metrics do not represent the performance of a system in its entirety. We follow this assessment and argue for a task-oriented view on the performance of a positioning approach that we introduce below.

Sensor data for the IPIN 2016 offline competition was collected by an actor walking along a predefined path as closely as possible, stopping at certain points to mark the ground truth (Torres-Sospedra et al. 2017b). For the offline competition at IPIN 2017, a few variations of carrying the device, e.g. by simulating a phone call, were added (Torres-Sospedra et al. 2018). In recent Microsoft Indoor Localization Challenges, only static locations were evaluated (Lymberopoulos and Liu 2017).

To our knowledge, none of the existing studies actually create a realistic scenario for the intended application context, i.e. in our specific case a navigation task in an area that is (at least partially) unfamiliar to the test person. The methodology presented by De La Osa et al. (2016) aims at real-life use cases, but still relies on a tester who can identify checkpoints on a predefined path.

3. Landmark-based Navigation

In the following section, we give an overview of the navigation system we use in the experiments that we discuss later in the paper. In particular, we describe how the graph model relates to areas and corresponding activities of indoor routes. Based on this description, we then introduce the particle filter implementation that forms the core of our indoor positioning approach.

3.1. *Data Model*

Fig. 2 visualises our data model that we call *indoor navigation graph*. It is used for computing routes, generating navigation instructions, and indoor positioning (see Müller

et al. 2017 for details). In order to be able to serve these three purposes, it formalises knowledge about activities pedestrians perform in the physical environment. All possible activities (see Sect. 1.2) are represented as edges in the graph, while the nodes represent locations where these activities can take place. In particular, source nodes of edges represent the location where an activity can start, and sink nodes are used for the location pedestrians are expected to reach after they have completed the activity successfully. The edge is typed by a unique activity out of the list above and indicates which kind of movement is to be expected next. In this way, we could potentially train classifiers for each edge type.

Furthermore, the graph contains nodes for objects pedestrians can observe in the environment (so-called landmarks). They are used to segment a route into areas. One navigation instruction is computed for each area. While navigating a user, the system has to keep track whether the user has completed the current activity as requested in the current navigation instruction. Following this approach, for locating the user it is sufficient to know that the user is close to a landmark (e.g. a door at the end of a corridor or a certain cloth shop in a shopping mall) as human users are capable to approach the landmark autonomously, i.e. a pedestrian can walk to a distant object without continuous technical assistance while a robot cannot. Therefore, the mean positioning error – while being definitely of interest for building autonomous systems that can navigate in indoor environments (e.g. robots for ambient assisted living) – for the implementation of many location-based services it is an inappropriate performance metric. Consequently, the precision of indoor positioning for pedestrian navigation systems should be measured in terms of the degree to which the current task-related activity is completed instead of in meters in a coordinate system that the user cannot even perceive (see e.g. Ohm, Müller, and Ludwig 2015).

At this point, we can state one major contribution of the present paper. We propose an approach to combine a graph-like representation of an environment with the minimally necessary metric information to correctly align the data computed by a SHS in the graph in order to assign the user's position to more easily perceivable objects in the environment which we call **areas** or **landmarks** depending on whether we refer to a part of a path the user should walk on or a relevant object the user can perceive in the environment. An example can be seen in Figure 4, which shows part of a corridor as an area consisting of a several adjacent edges. In Figure 5 a typical landmark is displayed: the billboard shown on the map is also referred to in the navigation instruction. This approach for an indoor navigation system shares similarities with other work (in particular Link et al. 2013). As a new contribution, we introduce the concepts of **areas** and the **area match score** that link indoor positioning based on SHS with landmark-based navigation (see Sect. 3.2). By doing this, we relax performance requirements for positioning algorithms as we no longer need to optimise the metric errors at any time of the navigation process. Instead, it is sufficient to identify the correct area a user is currently walking on: given the current area, the system can generate a navigation instruction that incorporates a landmark easily perceivable from the estimated position of the user.

In a similar fashion, Pulkkinen and Verwijnen (2015) remark that for certain applications it is more important to reliably identify the room where a person or object is located than to know the exact position. They therefore introduce zone/cell accuracy, operationalised e.g. as the classification error. Our approach, in contrast, does not rely on predefined, static zones derived from the building geometry, but generates its areas dynamically depending on the planned route and the available landmarks. In other words, our areas are not necessarily tied to the building layout but rather temporarily
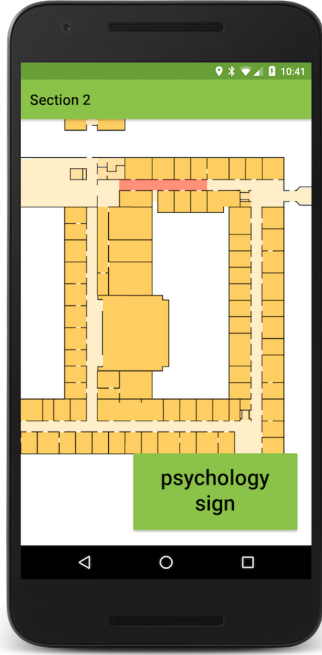
Figure 4. User interface of the data collection app for the initial study. The current area is highlighted in light red.
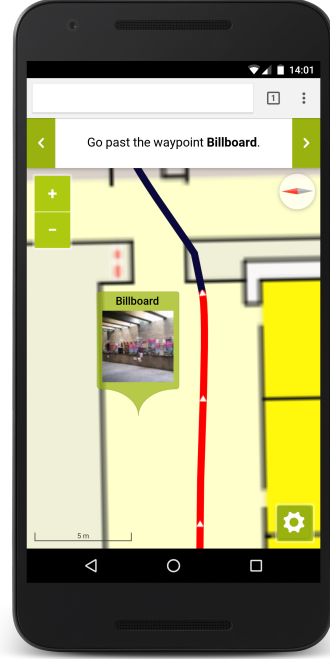
Figure 5. Example of a landmark-based navigation instruction.

partition the space in a way that relates to the current navigation task.

As stated above, for landmark-based navigation it is crucial to give navigation instructions at the right time in order not to confuse users and to guarantee good usability as well as reaching the destination. A navigation instruction is given at the right time if it does not refer to any landmark that is not yet visible from the user's current position or that the user has already passed. As a consequence, for indoor positioning the main difference to zone accuracy lies in the fact that the positioning algorithm has to determine whether the user is *close* to the landmark and can switch to the next instruction. Mostly, this task is much easier than continuously determining the exact position.

For the effort to create indoor navigation graphs, we note that it is impossible to extract all the information about decision points and likely taken paths algorithmically, especially when working with imprecise or outdated plans. The data we need goes beyond what is contained in automatically generated topological building models (e.g. Hilsenbeck et al. 2014). While there is no doubt that the automatic generation of models for indoor environments will see much progress in the future – in particular due to the application of modern machine learning techniques –, this issue is not in the focus of our work. Instead, by entering the information intellectually and – if needed – on-site, we can achieve a higher similarity between model geometry and actual trajectory. In summary, the system relies on a single data model for routing, instruction generation, and localisation that allows us to minimise the effort needed for map creation and maintenance.

3.2. *Graph-based Localisation*

In order to give the correct instructions at any time, the user's relative position towards landmarks referred to in navigation instructions needs to be known. Our indoor positioning algorithm computes this position by mapping sensor data to areas in the indoor navigation graph. For this mapping, we implemented a recursive stochastic filter that after each measurement assigns a probability to each area proportional to the likelihood of the user to currently walk on a certain area.

The filter is implemented as a particle filter (see Thrun, Burgard, and Fox 2005). This family of algorithms represents a probability distribution by means of a representative sample, a set of so called particles. Around the expected position the number of particles is high while elsewhere it is low according to the small probability mass. Using a sampling and resampling strategy the set of particles is updated after each measurement in order incorporate the new information (see Thrun, Burgard, and Fox 2005 for details): far away from the expected position the particles diminish while new ones are generated for positions with high probability mass. Our implementation can also incorporate input from multiple sensors such as detected steps or WiFi signals in order to implement the advocated hybrid approach for indoor positioning. Furthermore, information contained in the indoor navigation graph stabilises and corrects the position estimates as many constraints for the user's current location can be derived from the graph structure (in particular invalid positions receive probability zero while in standard SHS approaches the same locations are possible positions).

The implementation does not rely on a particular sensor technology or SHS algorithm, but only assumes to receive vectors that quantify the step length and direction of a pedestrian's movement. In fact, we do not perform any low-level sensor fusion or step detection based on raw sensor values. In order to investigate the influence of the precision of the SHS on our approach, we compared two different algorithms:

- motionDNA by Navisens
  The motionDNA SDK by Navisens is a well-known commercial state-of-the-art motion tracking solution. According to the company's website[2], it relies on inertial sensors only and does not need any external infrastructure to operate. The sensor readings are updated with a rate of 24 Hz on our test device and include a variety of information such as the user's activity and the device orientation and position. For this study, only the position information (relative to the initial position, measured in meters in X and Y direction) is used.
- Android's built-in sensors
  On recent devices, the Android framework gives access to many sensors that can be used for motion tracking. The step detection sensor tells us whenever a step occurs. Since it does not provide a step length, we initially use a fixed length and later use the particle filter to adapt to the user. Under the assumption that the user orients the smartphone roughly in his walking direction, the average orientation during the step as provided by the rotation vector sensor is used as step direction. Though this may seem like a strong assumption, we show empirical evidence later on that in the targeted scenario pedestrians tend to carry their device in exactly that way most of the time.

In Fig. 6 the data computed by each of both algorithms for a single walk on the test route of the initial study is plotted into the map of the building. Many position
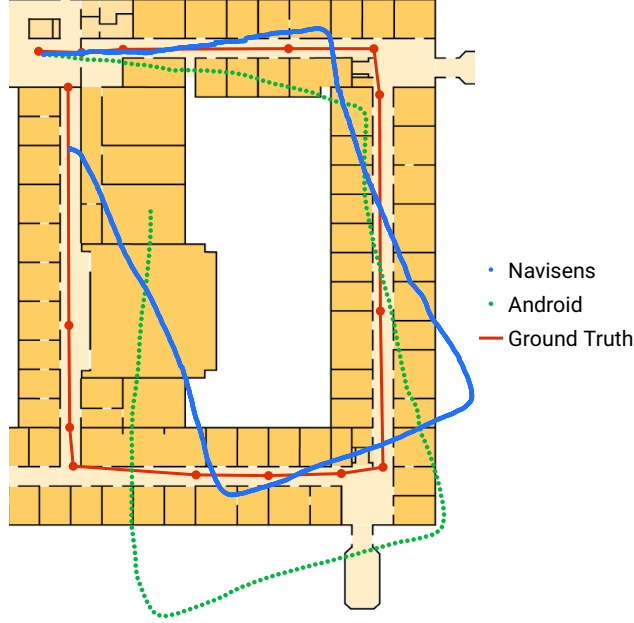

Figure 6. The trajectories of both motionDNA (blue) and raw step data (green) of a typical walk along the test route of the initial study. Ground truth (starting top left, then clockwise) and the boundaries of the areas defined for the evaluation experiment are drawn in red.

estimates are far off the route. This observation illustrates that information about the environment is indispensable for the position estimates to be used in an indoor navigation system.

In order to map SHS estimates to the indoor navigation graph, we apply the described particle filter. Initially, the probability is distributed uniformly over all edges. If information about the start position is available, e. g. during a navigation task, we use a normal distribution centred around the start node instead to sample the initial particle set.

Whenever a step is detected, a Gaussian naive Bayesian classifier updates the probability distribution for the edges starting in the current node. The update takes the motion model for the user (i.e. the distance and direction of the detected step) and the orientation of the considered edges into account. The probability of the user to walk on an edge increases if this edge is parallel to the direction detected by the SHS. The increment for an edge not parallel to the detected direction decreases proportionally to the angle between the direction and the orientation of the edge.

Unlike other approaches, the algorithm does not immediately select the edge with the highest probability as the current position estimate. Instead, it updates the set of particles each of which represents a different hypothesis for the user's current position. A similar approach has been successfully applied to localisation in robotics (see Thrun, Burgard, and Fox 2005) and allows to

- account for noise in the SHS data, which may stem from the rotation vector sensor (or rather the underlying magnetic compass) or the way the device is held in the hand,
- account for differences in step length while a person is walking, and
- account for different step lengths of different users.

More formally, each particle's state is defined by the vector $\{n_t, d_t, e_t\}$, where n_t de-

notes the starting node at time t, d_t the distance walked since leaving the node, and e_t a discrete probability distribution for the edges adjacent to the node. On every step, the state is updated according to

$$\{n_t, d_t, e_t\} \sim p(n_t, d_t, e_t | n_{t-1}, d_{t-1}, e_{t-1}, z_{\theta,t}, z_{l,t}, G), \tag{1}$$

where $z_{\theta,t}$ and $z_{l,t}$ are the measured step direction and length, and G is a indoor navigation graph. Applying the procedure detailed in Hilsenbeck et al. (2014), the update rule can be decomposed to its independent parts. The noisy step length measurement with the empirically determined variance σ_l^2 is modelled by

$$l_t \sim p(l_t | z_{l,t}) \sim \mathcal{N}(z_{l,t}, \sigma_l^2), \tag{2}$$

leading to the updated cumulative step distance of

$$d_t \sim d_{t-1} + l_t. \tag{3}$$

Similarly, the step direction is updated by

$$\theta_t \sim p(\theta_t | z_{\theta,t}) \sim \mathcal{N}(z_{\theta,t}, \sigma_\theta^2) \tag{4}$$

and subsequently used to determine the new edge distribution:

$$e_t^i \sim p(e_t^i | e_{t-1}^i, \theta_t, G) \sim \mathcal{N}(\Delta\theta_t^i, \sigma_e^2) * e_{t-1}^i \tag{5}$$

Here, e_t^i denotes the probability of the user to currently walk on the i-th edge adjacent to the current node and $\Delta\theta_t^i$ the angle difference between the step and the i-th edge. Finally, the decision whether the user has completed an edge and moved to the next is formalised as:

$$n_t \sim \begin{cases} \text{no} & \text{if } d_t \leq \text{length}(e) \wedge e = \text{argmax}_i(e_t) \\ \text{yes} & \text{if } d_t > \text{length}(e) \wedge e = \text{argmax}_i(e_t), \end{cases} \tag{6}$$

i. e. whenever d_t exceeds the length of the currently most probable edge e. In this case the starting node has to be updated: n_t is set to the sink node of the previous edge and d_t is reset to zero. Since the walked distance usually does not align exactly with the edge length, the difference is added to the position estimate and the step bias is reinitialised to $\mathcal{N}(z_{l,t}, \sigma_l^2)$ as the prior distribution for the new current edge e_t.

After the update step, the particle importance weights are distributed according to the non-normalised probability of the most probable of all adjacent edges:

$$\omega_t = \omega_{t-1} * p(z_{l,t}, z_{\theta,t} | n_t) \sim \omega_{t-1} * \max_i(e_t) \tag{7}$$

Finally, stochastic universal sampling is performed, which guarantees low variance and a representation of the samples in the new particle distribution that is proportional to their importance weights (see Thrun, Burgard, and Fox 2005 for further details).

In order to estimate the user's position, the expected value of the particle distribution is calculated. From there, the closest point that is located on either an edge or a node of the graph is computed as the final position estimate. This snap to the indoor

navigation graph ensures that the position estimate is a location that is accessible to the user and – differently to the pure SHS algorithms – prevents the positioning algorithm to assume impossible movements, e.g. through walls.

4. Task-Oriented Evaluation

As stated above, a central objective of our work is to observe the activities of users in the context of navigation tasks. As these activities correspond to certain traces of movements that can be observed by state-of-the-art sensors, we can exploit this dualism for activity recognition. In practice however, as we also noted above, the state of the art of tracking movements is still imperfect. One reason for this – and therefore a chance for improving the accuracy – may be the fact that currently for indoor positioning all existing approaches use one single algorithm even if the physical models underlying the different kinds of movements actually suggest that for each kind of movement a particularly suited algorithm could be appropriate. While this observation is obvious for *walking upstairs* compared to *walking straight ahead* in a plain space, in indoor positioning a distinction has never been made e.g. between classifying raw data into *walking straight ahead* and *turning right*. However, as research in biology and medical engineering suggests (see Novak et al. 2014; Nandikolla et al. 2017; Hase and Stein 1999; Sreenivasa et al. 2008), it could well be worth to train a different classifier for each type of movement.

As a first step in this direction, we wanted to analyse realistic data whether positioning errors of a state of the art-SHS-approach can be better explained if one knows the expected type of movement. It could, for example, to true that on average SHS-approaches perform well if users walk straight ahead, but could fail quite often to detect users to turn aside.

Therefore, in order to explain the limits of PDR-based indoor positioning systems, we have conducted three studies so far. In this section, we will describe each of them in detail and reveal typical situation in which the accuracy of indoor positioning is low.

The first study introduces the concept of task-oriented evaluation and the area match score, while in the second experiment we collected and analysed data under realistic conditions for everyday navigation scenarios. In the third study, we collected and evaluated new data, taking the lessons learned from the previous studies into account.

4.1. *Study 1: Applying the Area Match Score*

The purpose of the initial study is to demonstrate the benefits of a task-oriented performance metric that shows how well the correct area on a route can be determined by the localisation system, which – as noted above – is a requirement for correct navigation instructions and for successful navigation in general. From the pedestrian's perspective, the concept of area-wise navigation instructions and the area match as an estimate of the pedestrian's performance result in minimising the cognitive workload necessary for aided wayfinding.

The second purpose of this study is to identify contexts and situations during a navigation task that point to problematic areas and to propose ways to mitigate the revealed issues.
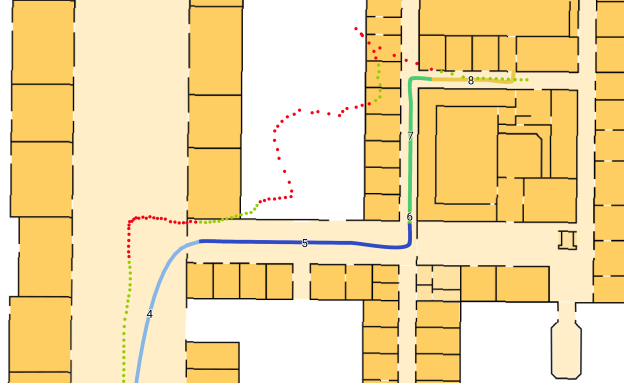
Figure 7. Detailed view of the results for a single walk on the second half of route 1 of study 3. Green dots denote successful position estimates; red dots represent estimates that are located in the wrong area or too far away from the route.

4.1.1. Setup

Based on the mathematical model introduced in Sect. 3.2, after each sensor data update t our indoor positioning algorithm predicts the user's current position p_t. If and only if the distance between p_t and the spline interpolating all nodes of the current area is sufficiently small, p_t matches this area (see Fig. 7). The user is now assumed to execute the activity corresponding to the calculated area. For this choice to be perfect, i.e. to be always correct, users have to walk exactly on the spline, and the sensor data has to precisely observe this movement.

To collect data for an evaluation of the implemented indoor positioning algorithm, we conducted an empirical study in the ground floor in an university building. There, we defined a test route spanning 182 meters. The route leads through 4 corridors in a rectangular shape. Three of the corners are modelled as small foyers (see Fig. 2). The only obstacles on the route are several glass doors that had to be passed in order to reach the destination of the route.

The route was segmented into areas as described in Sect. 1.3. Their boundaries were set at positions where semantically relevant objects – i.e. salient landmarks – are located. For determining salient landmarks along the route, we followed the approach described in Kattenbeck (2016): 19 persons rated 32 objects in the test area regarding different aspects of their salience. We selected the objects with the highest predicted overall salience as landmarks for the navigation instructions in our experiment. These landmarks included e.g. a glass cabinet, a wall painting, a bench and a sign for the department of psychology. Additionally, architectural features such as the aforementioned glass doors or the beginning and end of foyers were used to segment the route into areas. For each area, we formulated a navigation instruction that should explain to the test persons how to proceed the route. Finally, the route consisted of fourteen areas of varying size (see Fig. 6). The main factor that influences the size of the areas is the visibility of the landmark at their end: some can be referenced unambiguously from further away, while for others one has to be closer, thus causing smaller areas.

4.1.2. Procedure

Acquisition of Positioning Data

Starting from a defined position, 7 different persons who were familiar with the area

and the landmarks performed a total of 15 walks along the test route. Data collection took place over the course of several days, with an LG Nexus 5X running Android 7.1.2 as the test device. Before each test run, the compass was calibrated and its proper functionality was verified. During the experiment, the phone was held in the hand in front of the body, pointing in the direction the person was heading toward.

For data collection, a custom Android application was developed. It is able to capture data from various sensors of the device:

- Steps detected by the built-in Android step detection sensor.
- Orientation data from Android's rotation vector sensor, which in turn fuses magnetometer, accelerometer and gyroscope readings.
- Data from Navisens' motionDNA SDK. First and foremost, this includes the relative position, but also heading direction, orientation of the device, as well as detected user activity.
- The signal strength of nearby WiFi access points (not used in this study).
- A video recording of the device's back-facing camera, capturing the test person's feet and the area immediately in front of them.

The app's user interface consists of a map of the test area and a single button that allows the user to start the test run. At the beginning of each test run, the indoor positioning system is initialised on the starting node of the test route. Next, the first area to traverse is highlighted on the map and the corresponding navigation instruction is displayed. Each time a test person reaches the landmark related to the instruction, he or she has to press the button in order to set the ground truth for the transition between two adjacent areas. Finally, the interface is updated with information for the next area on the route. This procedure is repeated for each area of the chosen route.

Validation of the Collected Data

With this experimental setup, we collected sensor data for the test route and a ground truth labelled by experts in a single run of the experiment. As all test persons were instructed before the experiments how to label the ground truth, the data sample is valid for mathematical analyses.

In order to verify whether the collected samples were representative for average persons walking straight ahead, several gait characteristics were calculated:

- The mean gait speed during a walk can easily be determined by the quotient of route length and the time needed to complete the route, measured by the difference of timestamps between last and first step. The result is a mean speed of 1.30 m/s ($SD = 0.14$ m/s), which is well within the margin reported by Bohannon and Williams Andrews (2011).
- In order to calculate the step length, the steps are counted manually for each walk by means of the recorded video, revealing that Android's step detector misses about 5.8% of steps on average.
- The mean step length amounts to 0.73 m ($SD = 0.077$ m), which is classified as fast gait according to the study from Oberg, Karsznia, and Oberg (1993). This can be explained by the fact that the test persons knew the area and the route very well.

In summary, the collected data is representative for the activities of humans during the pedestrian navigation processes we want to analyse.

4.1.3. Analysis of the Collected Data

The analysis of the raw data shows – quite expectedly – that the error quickly accumulates, leading to a high mean location error of 11.5 m (Android sensors) respectively 12.0 m (motionDNA). Figure 6 shows the trajectories of a typical walk. Navisens' motionDNA often struggles with substantial drift towards the left early on, but otherwise manages to track the overall shape quite well. The version relying on the Android step counter usually shows drifts in different directions throughout the walk due to the lack of correction. Additionally, the reported distances differ between the tracking methods: motionDNA's paths are usually shorter ($M = 174.1$ m, $SD = 15.37$ m), Android's longer ($M = 189.4$ m, $SD = 16.49$ m) than the ground truth of 182.0 meters.

Before the motionDNA data could be used as input to the particle filter, some preprocessing was inevitable: Since the update frequency of about 20Hz was rather high (about an order of magnitude higher than the step frequency), the data was split in batches of ten measurements that were treated as a single step. In two of the 15 recorded walks, the relative location reported by motionDNA unexpectedly was set back to the starting point of the route. Therefore, the area in which the reset occurred was eliminated from the data set.

4.1.4. Results

Since it was not feasible to run both Navisens' and our indoor localisation implementation at the same time on one device, we processed the collected data in an offline simulation of our indoor positioning algorithm.

In order to extend the data set, sensor data from each actual walk was used multiple times with different initial random seeds for the particle filter, resulting in 10 iterations for motionDNA and the Android sensors, respectively.

The extended data set was used for the evaluation of the implemented algorithm. In the remainder of this section, we present our evaluation results and discuss their impact on the appropriateness of the proposed area match score for localising users during indoor navigation.

Accuracy Metrics for the Sensor Data

Figure 8 shows a comparison of the two motion tracking solutions regarding their positioning accuracy, i. e. the distance from estimated position to ground truth, after their raw data has been processed by the particle filter.

The mean and median error of motionDNA amount to 7.02 and 4.28 meters respectively, while the Android sensors lead to an accuracy of 4.39 (mean) and 2.60 (median) meters. This performance gap is likely caused by two factors:

- The drift at the beginning that motionDNA often suffers from is propagated throughout the whole walk, causing a mismatch between step directions and the graph edges.
- Open spaces at the ends of the corridors allow for some overshooting, which benefits the approach using Android sensors and its slightly longer steps. The too short distance reported by motionDNA however can often not be compensated by the particle filter.

Analysis of the Area Match Score

On average, 60.6% (motionDNA) resp. 74.7% (Android sensors) of position updates are assigned to the correct area. Figure 9 visualises the area match score for each area
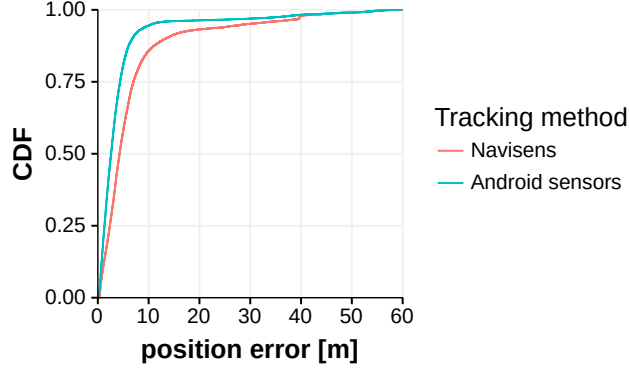
Figure 8. Empirical cumulative distribution function showing the accuracy with the two different motion tracking methods.

on the route. Obviously, the choice of the SHS influences the overall performance of our positioning algorithm. It cannot repair arbitrary errors of the SHS as positions too far away from any edge and directions very different from the orientation of the edges nearby the user's current position decrease the probability of the particles for these edges significantly (see Eq. 5).
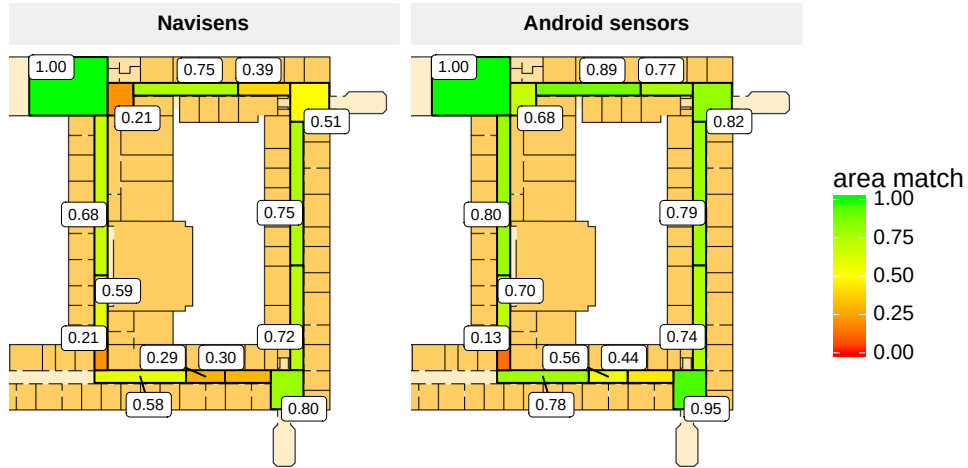

Figure 9. Area match scores for the two motion tracking methods.

As Fig. 10 illustrates, the area match score and the positioning error are inversely correlated ($r(58)$ = -0.87, $p < 0.05$). From these observations we conclude that in order to support indoor navigation effectively any indoor positioning needs to be able to reliably estimate a user's relative movements. While in this study we only analysed walking, this observation in a more general setting equally applies to other kinds of movement (e.g. climbing stairs, taking an elevator, etc.).

Influence of the Navigation Graph on the Area Match Score

While from the preceding analysis we learn the lesson that the area match score's precision depends on the quality of the step detection, in the following we identify
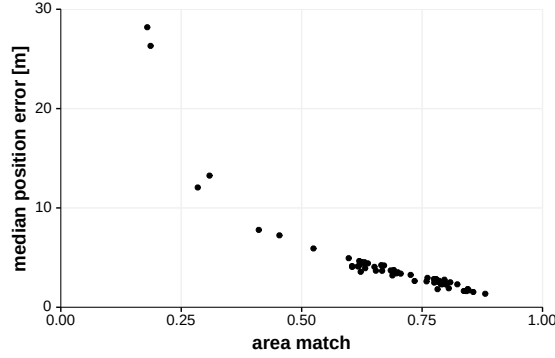
18

Figure 10. The area match score correlates inversely with the median positioning error. Each point represents the mean of 10 simulated runs for one actual walk.

other sources for area match errors.

The first source is the indoor navigation graph. Its usage introduces artifacts for the actual movement of a person as it always has to be snapped on one of the edges – sometimes a very crude discretisation of the actually available degrees of freedom how to move.

While in corridors no problems may arise, Fig. 9 illustrates that in junctions and foyers, the area match score tends to decrease. The limits of *delay*, *magnetic bias*, and *step length* amount to the problem of *turn detection*. The sensors obviously are not fast enough to reliably reconstruct the pedestrian's motion – this is a clear hint for the need to develop classifiers for turns as part of future work in indoor positioning. Such a classifier should take the context (i.e. the activity users are expected to perform according to the current navigation instruction) into account.

By analysing the data and the indoor navigation graph, we revealed a second issue that influences the area match score: junctions of two corridors are often modelled using three edges only (see Fig. 11). In such situations, there is only a single correct edge that can be hypothesised as the current position. However, the particle resampling may fail when the SHS due to the limits mentioned above misses the user's turn or at least recognises it too late. In this case only few or even no particles are generated for the current edge while the majority of the particles hypothesises the user to continue to walk straight ahead. This phenomenon is particularly obvious for the junction on the bottom left in Fig. 9.

Contrarily, the three other larger foyers are represented by a densely connected net that enables the system to track almost arbitrary paths within these areas (see Fig. 2). In these areas, the area match score remains high.

This circumstance teaches us that indoor navigation graphs should not only model accessibility relations between locations in the modelled environment, but also approximate the geometry of the locations.

We tested this hypothesis by connecting the nodes adjacent to junctions with additional slanted edges as depicted in Fig. 11, the benefits of which are twofold: Firstly, it models more natural paths where the test person cuts the corner slightly; secondly, it allows for the compensation of step length differences since now multiple paths lead into the corridor that is branching off.

Using the new graph structure, we repeated the computation of the area match score. The result was not only an improvement in the area after the junction, but in all subsequent ones as well. In the small, but critical area immediately after the junction,
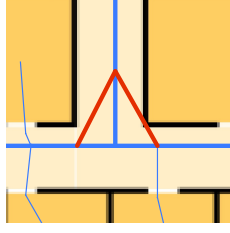
Figure 11. Closeup of the part of the graph that was changed. The edges drawn in red were added to stabilise the position estimation at this junction.
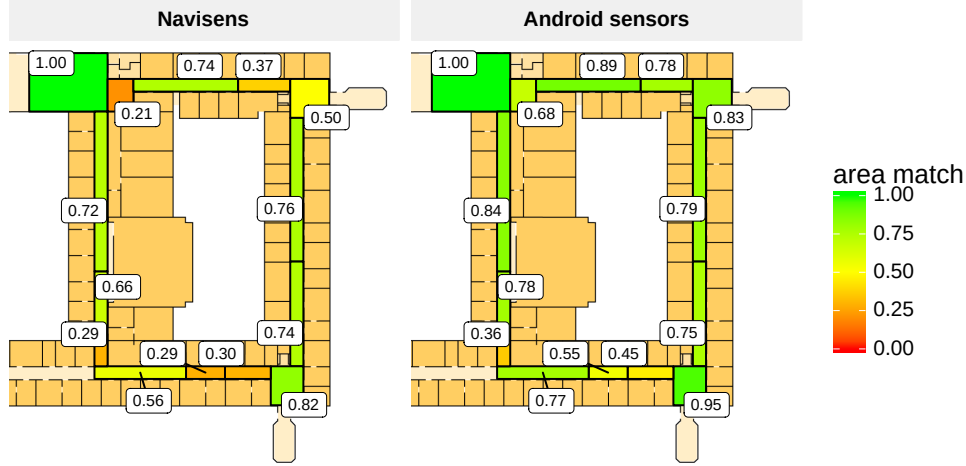

Figure 12. Area match scores for the two motion tracking methods (with additional edges in the lower left corner).

the area match score was improved by 38% (from 0.21 to 0.29) for motionDNA, and almost tripled (from 0.13 to 0.36) for the Android sensors. The improvement is even statistically significant for the remainder of the route after the change: a Wilcoxon rank sum test indicates that the area match score is greater for the graph model with additional edges ($Mdn = 0.78$) than for the original version ($Mdn = 0.73$), $W = 33942$, $n_1 = n_2 = 300$, $p < 0.05$.

For reference, the median position error when calculated for the whole route also decreased from 2.60 to 2.52 meters for the Android sensors, and from 4.28 to 4.11 meters for the version running with motionDNA.

We conclude that by applying a systematic methodology to design an indoor navigation graph, we can almost completely eliminate the negative influence of the discretisation of the physical environment into edges and thereby implement an approximate solution for a *turn detection* classifier working on sensor data. We will have to investigate whether combining both approaches will lead to a reliable turn detection that seems to be a challenge specific to indoor positioning.

A further note: the comparison of the two motion tracking solutions showed that the supposedly more sophisticated one does not outperform the built-in step counter when embedded in a more complex, non-metric approach to indoor positioning. Both suffer from a cold start problem and produce wrong estimates when indoor positioning starts. To Navisens' credit, we only used a small portion of motionDNA's capabilities and designed the experiment in a way that the Android sensors would have a reasonable

chance at competing, e. g. by restricting the device location and only using a single floor for the test route.

In the remainder of this analysis, we only discuss results obtained by applying best practices learned so far: we use the internal Android SHS on the indoor navigation graph with additional edges for junctions (as shown in Fig. 11).

Influence of Area Transitions on the Area Match Score

Dividing a route into areas as explained above introduces another artifact at the boundaries of adjacent areas. It may prove problematic that boundaries are strict while SHS is noisy. Therefore, measurements taken around boundaries may be randomly assigned to one of the areas and increase the area match error.

In particular, the smaller an area is, the higher the precision of the SHS has to be for the measurement to be matched to the correct area. Therefore, in order to eliminate the influence of this artifact on the area match score, it seems justified to smooth the boundaries, allowing positions up to 2.5 m (i. e. the median position error) away from the exact boundary still to count as a match.

By loosening the definition of an area match in this way, the score increases from 0.77 to 0.88 on average, almost cutting the remaining error in half. Considering only the middle part of each area, defined as those positions that are further than 2.5 m away from each of the area's boundaries, the area match score amounts to 0.87 (strict) respectively 0.91 (approximate). On the other hand, when looking at the boundaries themselves (i.e. the interval of ± 2.5 m around the boundary), the scores amount to 0.73 for the parts immediately after a segment change and 0.84 for the part at the end of each segment.

In summary, we conclude that the SHS position estimate tends to lag behind more often than it precedes the actual position.

4.2. *Study 2: Creating Realistic Conditions*

As detailed above, the initial study uses a controlled procedure for data collection, where an expert walks along a known path as precisely as possible, annotating the ground truth at predefined positions. This approach – also described by e.g. Torres-Sospedra et al. (2017b) – has quite a few benefits, e.g. a high degree of repeatability. Furthermore, it might be chosen simply for practical reasons since a few persons can quickly collect a large amount of data.

However, we argue that data captured in this way does not convey the full range of human activities during a real navigation task. Therefore, it is not possible to guarantee that the accuracy achieved with such 'artificial' data sets can be reproduced in real-world applications. As a consequence, it would be ideal to perform naturalistic studies. However, they are impractical for the purpose of evaluating indoor positioning systems since we lack a high precision approach that could automatically observe ground truth data automatically in naturalistic situations. Lately, there have been advances towards a more realistic setup, e.g. by simulating different ways of holding the device (Torres-Sospedra et al. 2018), but they still rely on an expert who knows the environment or the route in advance.

Therefore, based on our own experience and hints for best practice reported in the literature, we devised a new experimental setup in order to collect data in conditions as realistic as possible.
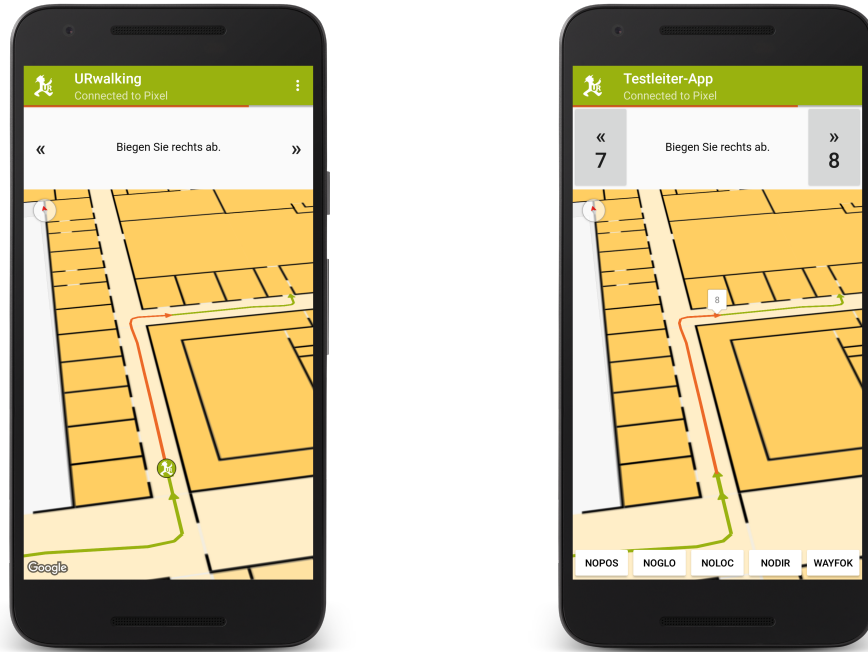
(a) App for the test person, displaying the route (current segment highlighted by colour), position indicator, and navigation instructions.

(b) App for the experimenter, with additional UI elements for ground truth annotation.

Figure 13. Screenshots of the data collection setup

4.2.1. Setup

In the new setup, the test person was given an actual navigation task that started at the current position and led to a predefined room somewhere on the university campus. A prototype of our navigation app for Android showed the route and provided navigation instructions, while in the background all relevant sensor data was logged. Navisens' motionDNA was not used this time since it did not show any significant improvement over the integrated sensors in the initial study.

The experimenter walked close by and used a separate application that allowed him to annotate the ground truth position and to control the test person's navigation app remotely via Bluetooth (see Fig. 13). The test person's app was configured to not switch instructions automatically when a new area was reached according to the positioning system, but only when the corresponding signal was received from the experimenter. Test persons were encouraged to think-aloud, and both devices were recording the conversation in order to gain insights into the mental state of the test person.

The main advantage of this setup is that the test person can concentrate on the navigation process and behave as naturally as possible while still in a somewhat controlled experiment situation. In particular, he or she does not need to know or even stop at the points where the ground truth is annotated: only the experimenter needs to know the environment beforehand.
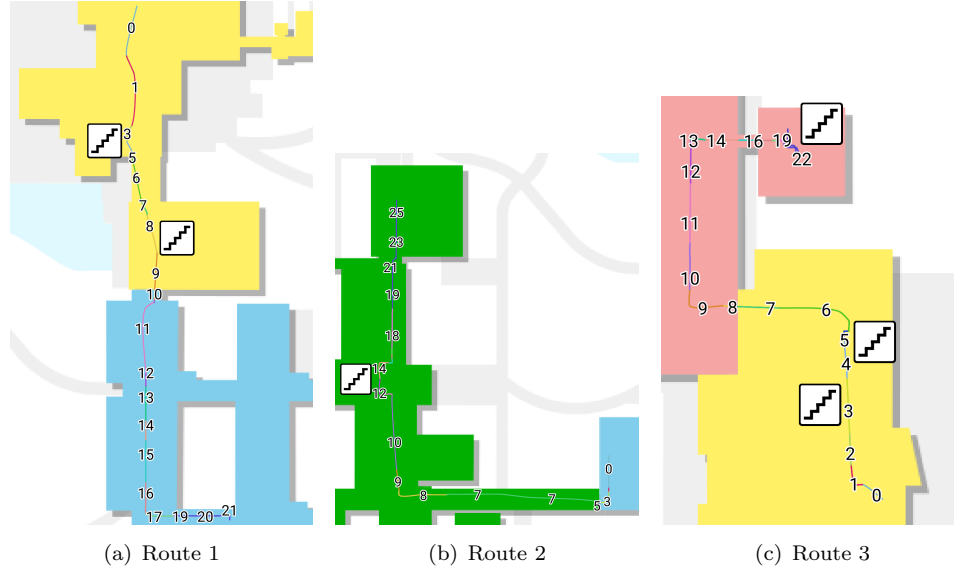
| (a) Route 1 | (b) Route 2 | (c) Route 3 |

Figure 14. The three different test routes used in the second study.

The experiments were conducted by students taking part in a course on HCI. Each group had to implement a different strategy to assist users e.g. when they had gotten lost, or while they had trouble finding the correct path. For the purpose of present paper, only the sensor data collected in the experiments under realistic conditions were used. The student groups acquired test persons who did not participate in the course on HCI.

The three long test routes (Fig. 14) ranging from 251 to 282 meters led across several buildings on the campus and were designed to cause situations where the test person would need assistance. Each of them contained at least one floor transition and, in addition, featured different obstacles and difficulties such as (revolving) doors, indoor/outdoor transitions, and highly frequented areas. So, we consider the routes as representative for everyday navigation situations as they contain all activities that humans eventually perform during navigation tasks.

Each group was provided with a different Google Pixel phone used by the test person. All in all, there were 114 runs logged by 38 test persons across 6 groups, resulting in roughly 91.2 km or 733 minutes of log data.

4.2.2. Results and Lessons Learned

The experiments provided valuable data about the user behaviour during a realistic navigation task: As expected, test persons occasionally slow down, stop to look around or need to re-orient themselves. During this time, the sensors sometimes still detect steps that of course influence the position estimate. These observations emphasise our point that controlled data collection by a professional leads to data that does not convey the full range of pedestrian movement. Systems trained on such data are likely to fail if applied to real-world scenarios and in this way exposed to previously unseen data.

One negative consequence of the new experiment design was the increased workload for the experimenters: They had to interact with the test person, annotate the ground truth, and determine which one of several assistance strategies to trigger whenever the test person exposed a certain behaviour, e.g. hesitated for a while or was unable
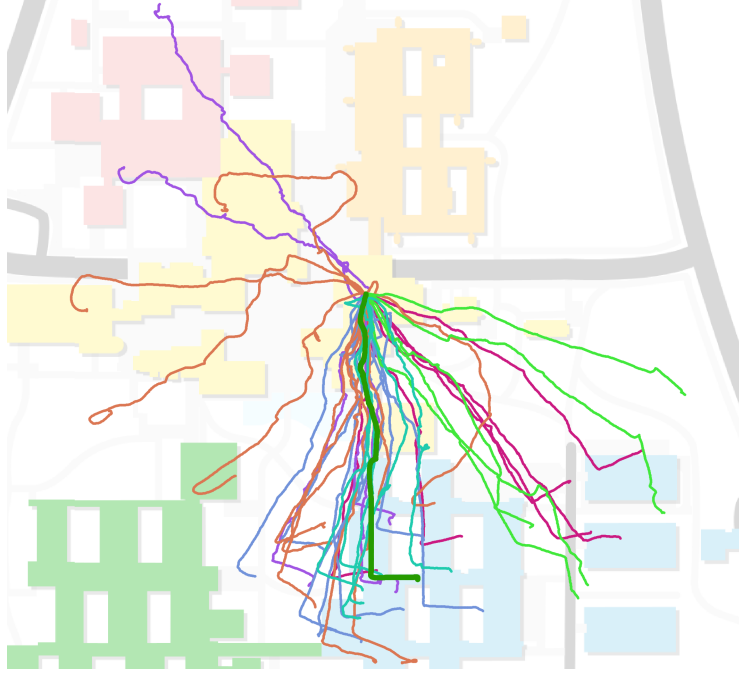
Figure 15. Raw step data from the first test route (dark green). Colour-coded by test device.

to understand the instruction. Therefore, the ground truth was not always annotated correctly, as apparent from areas for which less steps were logged than actually needed to completely traverse the area.

For this reason, we decided not to calculate the area match score or other performance metrics from this data. Instead, we focused our evaluation on the analysis of the sensor accuracy as a basis for future indoor-tracking experiments.

Since the compass was not calibrated before each test run started (we omitted the calibration in order to simulate realistic conditions), the estimation of the step direction proved to be more difficult: as can be seen in Figure 15 on the example of the first route, the measured directions spread around the true direction, with a mean error between 30 and 39 degrees and a standard deviation of 34 to 58 degrees when considering the routes as a whole. Also apparent is a systematic bias dependent on the test device, visualised by the different colours used to display the paths.

Looking at the individual trajectories, we often observe that turns are not registered correctly by the orientation sensor, in particular apparent by the ubiquitous 90-degree angles inside buildings. Immediately after these turns, the direction sometimes seemingly tries to self-correct, resulting in inconsistent drift in one or the other direction. This poses a big problem for the localisation system and can only to some degree be corrected by our graph model.

We note that in the second study we are again confronted with the problem of *turn detection*: limits of the smartphone's sensors and the SHS lead to a repeated failure to detect the activities of test persons reliably. In our view, the observed sensor behaviour may explain to a large extent why accuracy rates in typical IPIN experiments do rarely increase: it seems that a single classification algorithm does not reliably solve the indoor positioning problem. Instead, depending on the current task context and the activity that a user is expected to perform appropriate classifiers have to be selected.

In order to handle the often erroneous initial direction, we revised our implemen-

tation and approached the problem algorithmically: The particle state is extended by an element that contains the current heading bias, i.e. the angle that has to be added to the measured step direction to obtain the true heading. Initially, the heading bias is distributed randomly from 0 to 360 degrees for all particles. The particle filter then during the first few steps can adjust the expected value for the heading bias automatically if the user follows the route. Furthermore, the particle filter can alter the bias slightly on every resampling step to account for changing environment conditions. In a similar manner, heading drift compensation is implemented by adding a randomly initialised amount of drift for each particle. However, these measures cannot completely compensate erroneous SHS data – adding a new point to our argument that the limits of SHS may be eliminated by activity-specific classifiers resulting in increased area match scores.

4.3. *Study 3: Evaluating on a Precise and Realistic Data Set*

In the third study, we combined the evaluation methodology from the first study and the experimental setup from the second study. The goal was to collect more step and orientation data and a reliable ground truth in order to be able to apply our task-oriented metric onto a data set captured under realistic conditions. Based on the lessons learned, we slightly revised our setup:

4.3.1. Setup

In study 2 the test system calculated the estimated position, then searched the nearest point on the route and used it to display the user position. This procedure should ensure that the test person would not get confused by impossible trajectories, e.g. through walls. Some confusion, however, arose when the compass behaved erratically, sometimes causing the position indicator to move in the wrong direction. To avoid such confusions in the new study, one option to mitigate this issue was to stop displaying the position altogether, although we knew from feedback that test persons prefer their position to be displayed on the map. For the new study, we therefore opted to display the expected position and update it after each step. In this way we achieved the familiar feedback without the risk of influencing the test person unintentionally.

As in study 2, the evaluation took place on three different test routes (Fig. 16). Since the focus was solely on positioning again, they were not specifically tailored to trigger situations where exceptional assistance was needed. Still, each of them contained many features commonly seen in office buildings, e.g. narrow corridors, open spaces, staircases, doors, different kinds of flooring, and crowded areas. With lengths between 151 and 190 meters, the routes were shorter this time and confined to a single university building. The route segments and corresponding navigation instructions were generated automatically by the routing system once again.

In order to guarantee an exact and reliable ground truth, the same person was acting as experimenter for all test runs. As noted above, a simulated position was displayed on each step to avoid influencing the test person by wrong position estimates. Just like in a real-world scenario, no compass calibration was done prior to the experiments, and the test persons were not instructed to hold the device in any specific way. Both experimenter and test person were using Google Pixel phones running Android 8.1. For all three routes, 12 runs were logged, yielding 110 minutes or about 6.0 km of labelled data in total.

| (a) Route 1 | (b) Route 2 | (c) Route 3 |

Figure 16. The three different test routes used in the third study.

4.3.2. Results

Though test persons were free to carry the device in any way they wanted, they almost at all times opted to hold it in front of their body, pointing roughly towards the walking direction. This observation can be explained by the fact that the test persons had to read the navigation instructions and search the next intermediate destination on the map presented on the smartphone's display. In turn, this means that the device direction is acceptable as a proxy for the walking direction during pedestrian navigation.

Since our indoor navigation graph does not explicitly model the shape or the amount of individual stairs in staircases, we ignore them for the purpose of this evaluation. Instead, we reset the position after the staircases in route 2 and 3 and do not use the areas in question for the calculation of any performance metrics. As we'll discuss below, we consider the detection of climbing stairs a problem that has to be solved separately.

Analysis of the Raw Data

Similar to the previous study, due to compass bias and drift the raw step trajectories do hardly resemble the path that was actually taken, in some cases making it even hard to map them to the actually expected locations in the environment intellectually (see Fig. 17).

However, most of the measured walking patterns tend to look quite similar on visual inspection, with some areas in which the observed bias is consistent across all test persons, e.g. at the start of route 3. If this bias was constant throughout a complete test run, the indoor positioning system could compensate for it. However, we observe from the data that the bias actually changes while test persons are walking, e.g. after the staircase in route 3.

In addition to these location-specific variations, the step logs exhibit different characteristics depending on when the data was collected: Towards the end of route 2, for example, one can easily distinguish between two clusters of six runs each (see Fig. 17). The data for each cluster was collected on the same day, while there are a few days between the test runs of the first and the second cluster.

As in both previous studies, the problem of *turn detection* can be observed also in the new data set. In particular, 90-degree turns are rarely identified correctly by
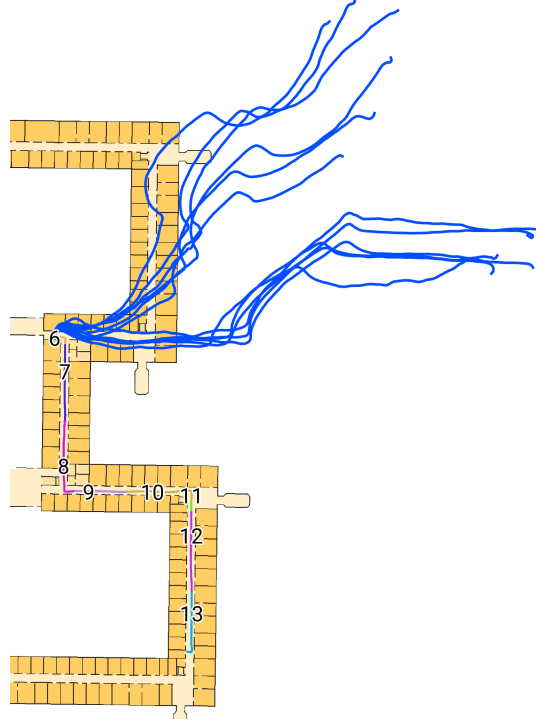
Figure 17. Detailed view of the raw step trajectories on route 2 (second floor only). Note the large discrepancy between actual route and captured data, as well as the different characteristic of the six runs performed at a later date (pointing roughly northeast).

the SHS. Techniques such as (improved) Heuristic Drift Elimination (Dez et al. 2016; Deng et al. 2018) that are commonly used to correct the errors introduced by compass or gyroscope drift are not able to handle the kinds of inconsistent deviations seen in our data. They would have to be configured to tolerate a high amount of drift, in turn making it even harder to detect where an actual turn took place.

For this study, the ground truth is available. The good news is that with the improvements of the indoor positioning algorithm and the assumptions mentioned above – variable compass bias that is learned by the first 15 steps, position reset after staircases –, the average area match score increases: our baseline system can now assign SHS position updates to the correct area in 73% of all cases. Without the assumption of the correct initial direction, the score suffers from the sometimes wildly inaccurate direction and decreases to 0.59. Due to the different experimental setup and the changes in the system, these figures are not directly comparable to those of the initial study. However, it again helps in exploring the limits of SHS for indoor positioning in navigation contexts: we can identify typical situations in which the area match score is low. Furthermore, we can provide empirical evidence that area match errors can be rooted back to wrong position estimates that, in turn, are caused by (some of) the limits of SHS we have identified already. As the remainder of this section will illustrate, we can even increase the area match score further by taking simple measures that simulate improved approaches for activity recognition:

Influence of Position Resets along the Route

Given the experience from study 1 that turns after long straight segments pose problems, for the data analysis we take a closer look at the individual runs. One area that

27

proved especially difficult is the branch on route 1 that leads away from a large open space through a door (see Fig. 7, end of segment 4), not unlike the situation in the initial study. Using this example, we explore to what extent the position estimate can be improved by precise position resets on a varying amount of points along the route.

With no position reset at all, the overall area match score across all test persons for route 1 amounts to 0.63. Resetting the position whenever a door is passed as well as at the end of a stairway results in a dramatic improvement, yielding a score of 0.90. With just two resets at locations that prove especially problematic, a score of 0.89 is achieved, and further reduction to a single point of re-initialisation at the aforementioned junction still yields an area match score of 0.83. In our opinion, these figures can compete with other approaches to indoor positioning mentioned in the introduction and have the potential to be improved even further.

Route 2 is characterised less by doors that have to be passed (most of them were always open during data collection and therefore would not lead to any detectable motion), than by its 90-degree corners. If we take the information of when these corners are passed into account, the overall area match score for route 2 improves from a baseline of 0.75 to 0.85. Reducing the amount of resets to the two with the most impact and then just a single one yields scores of 0.83 and 0.77, respectively. Finally, route 3 contains only one door where it is possible to initiate a position reset. However, due to its location near the end of the route this does not improve the overall score of 0.81 significantly.

These figures illustrate the importance of detecting activities such as *walking upstairs*, *walking downstairs*, *taking an elevator*, and *taking an escalator* and deduce the user's current position from the observed activity instead of detecting subsequent positions which may be quite noisy and infer the current activity from geometric properties of the last few position updates.

Further, and less obvious, evidence for this claim comes from our trial to observe changes in the user's orientation relying on sensor data not considered so far:

Eliminating Errors Introduced by the Magnetic Field

Since the error introduced by the Rotation Vector sensor repeatedly proved problematic, we explored as an alternative the Android Game Rotation Vector sensor, which does not rely on the magnetic field (Google 2017). It therefore cannot provide an absolute direction, but shows a much higher fidelity of the relative changes in direction (see Fig. 18). In particular, 90-degree turns are immediately recognisable, and though there still exists some amount of drift, it can easily be compensated.

To confirm that the gyroscope on its own does not suffer from substantial bias, we calculate the bias stability (i.e. the minimum Allan Deviation), which for our test devices amounts to completely acceptable values between 0.00002 and 0.00004 rad/s.

Since our results show that the initial direction is unreliable even with the magnetic compass included (see Fig. 17), the drawback of not knowing the absolute heading does not matter too much in practice; in fact, the same mechanism that stabilises the direction can also be applied here.

The similarity of step trajectories and ideal route can be quantified by the smallest warp path calculated by the Dynamic Time Warp algorithm (Magdy et al. 2015). Averaged over all runs, the rotation invariant smallest warp path decreases from 4563.6 to 2555.0 when not relying on the magnetic field, confirming the visual analysis. Repeating the simulation with these step trajectories as input, the average area match scores improve to 0.80, 0.86 and 0.84 for routes 1-3 without any position reset except
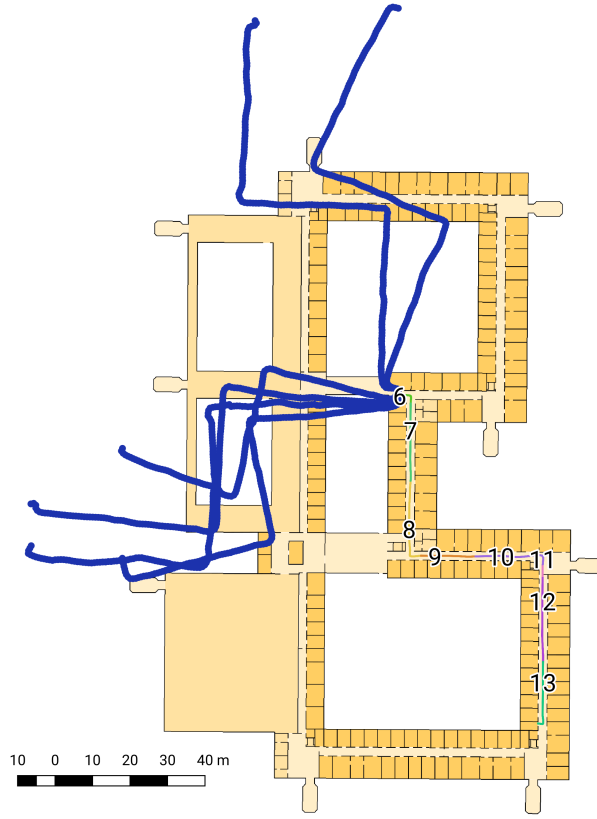
Figure 18. Detailed view of the step trajectories on route 2 (second floor only) without relying on the magnetic compass.

after staircases.

5. Discussion

The repeating theme in all three studies is the difficulty to track the user whenever the walking direction changes or one of the other identified limits of SHS negatively influences its position updates. In the initial study, areas immediately after 90-degree turns are amongst the ones with the lowest area match score, even after improving the graph model and thus allowing the filter to spread the particles over more edges that can better match the raw data (see Fig. 12). Creating more realistic conditions in the second study, e.g. by refraining from calibrating the compass, in many cases led to an erroneous initial direction. Furthermore, we observed that the angle of turns would often not get detected correctly and that the rotation vector would "bend back" during straight segments, posing a big problem when matching the trajectories to the indoor navigation graph. All these issues are present even more clearly in the third study, with an additional temporal influence becoming apparent.

The underlying cause, of course, is the limit of *magnetic bias* that is slow to react to changes and easily influenced by external magnetic fields. This effect should be minimised by the usage of Android's Rotation Vector sensor which fuses accelerometer, gyroscope and magnetic field sensor in order to detect orientation changes more

29

quickly. However, as our results in study 2 and 3 indicate, the sensor's capability to accurately determine the orientation is greatly reduced under realistic conditions. Under these circumstances, it is advisable to only rely on gyroscope and accelerometer (i.e. the Game Rotation Vector sensor), even though the ability to determine absolute heading information is then no longer present.

One major advantage of our graph-based filtering approach is the system's ability to correct the position back on the correct track even without external, absolute measurements (see Fig. 7). As we've seen however, and as it is common for dead reckoning systems in general, when at some point the error accumulates to a degree that is no longer correctable, there is a need to reset the position.

As detailed in the previous section, in many cases just one or two resets at strategically chosen decision points can greatly improve the positioning accuracy. How to determine these points based on a planned route and knowledge of the environment is subject of ongoing research.

One option traditionally used in hybrid systems to stabilise the position is signal-strength-based WiFi fingerprinting. For the purpose of resetting the position at certain exact locations, however, WiFi is not well suited since position updates are too infrequent and not sufficient in terms of precision. An alternative approach for this issue based on Bluetooth is presented by Hoffmann et al. (2017), but there's still no solution for the known drawbacks such as costly and time-consuming deployment and maintenance.

Therefore, as mentioned already in the discussions of the presented studies, we propose a hierarchical architecture of classifiers for different purposes. The high-level classifier uses raw data from the inertial sensor as well as others, e.g. the barometer and magnetometer, to create hypotheses about the current activity. The result from this classifier can be matched easily with the current context of the last navigation instruction in order to decide which of the created hypotheses are reasonable in the current context. The most probable among them should be the hypothesis that coincides with the expected activity – in this case, the user executes the navigation instruction correctly, while for all other hypotheses the user has to be warned about getting lost.

We simulated the effects of such an architecture by manually resetting the filter (see the discussion on the influence of position resets in section 4.3.2), yielding evidence that the strategy is reasonable.

A possible point of criticism against the proposed area match score is that it is averaged over the whole test route. That way, if localisation fails in a few small but important areas, the overall score can still be somewhat high even though localisation as a whole cannot be considered successful. This is correct, but the same holds true for other commonly used metrics such as the mean position error or the third quartile of the position error, which are also not influenced very much by a few bad estimates. Besides, any small malfunction or imprecise positioning as well as many other things such as badly worded instructions can lead to an unsuccessful navigation. The area match score as a single score cannot reflect all these possibilities, but can just give an overall impression on how the system works in most cases. For a more detailed analysis, it is necessary to directly look at the segments in question. To gain further insights, it might also be useful to report different metrics such as the minimum score and the standard deviation of the score, or to weigh the areas of a route in a different manner.

6.   Future Work and Conclusion

The main objective of our work was to explore the limits of PDR-based indoor positioning under realistic conditions. Before moving on to the challenges for future work in PDR research, we will recap the outcomes and findings of our studies so far:

- The area match score was introduced and validated as an approximate indoor positioning metric tailored to the needs of landmark-based navigation for pedestrians.
- Determining the walking direction based on built-in sensors proved difficult under realistic conditions, even when assuming a fixed device position. At least in the devices we tested, the magnetometer can be biased very easily by external influences, leading to early failures in position tracking. In such cases, e.g. turns cannot be detected correctly.
- Not relying on the magnetic compass altogether and instead inferring the initial direction from the information available from the map and/or the planned route led to a more accurate trajectory and a higher area match score.
- We demonstrated that long-term stability for PDR could also be achieved by position resets at certain decision points along the route, provided by a reliable complementary classifier.

With this in mind, we identified room for improvements in indoor positioning, coinciding with important research directions in indoor positioning in general and PDR in particular:

- **Context Model:** The graph model works best in narrow corridors with little degrees of freedom, but can be too restrictive in large open spaces. There, different factors are at play that have to be balanced: The model needs to accurately describe all paths that can be taken by a pedestrian, but cannot be too general or else it will not provide any stabilisation. Furthermore, the effort that goes into the manual creation as well as the computational effort resulting from a detailed model has to be considered. Hybrid models that incorporate open spaces into graphs as polygons (see e.g. Koivisto et al. 2015) provide a solution for the oftentimes heterogeneous nature of building layouts and should therefore be explored further in the context of PDR.
- **Motion Model:** We assumed the user to steadily walk ahead on the same floor in order to be able to reliably analyse the SHS data. However, analogous models have to be developed for other ways to move (in areas of other environmental type), such as taking stairwells, elevators, or escalators. According to the most recent position estimate an indoor positioning algorithm will have to decide at run-time which of the models to be used for the analysis of raw sensor data.

  As already addressed above, reliably determining the walking direction in unfavourable conditions (e.g. a fluctuating magnetic field) is still an open problem whose solution most likely needs to involve information from the map and from other sensors. While we observed that the users carried the device in front of their body at almost all times during navigation, a system that aims for near-perfect position tracking under all circumstances will have to account for other transportation modes as well. The same holds true if non-visual means of information transfer such as vibration or speech are used, where the user is no longer required to keep an eye on the smartphone screen to receive navigation updates. In our future work, we will investigate the influence of these alternative modali-

ties on the positioning accuracy.

- **Intelligent Sensor Fusion:** Finally, there is a lot of potential for optimisation in the way the various available sensors and classifiers are integrated with each other into a coherent system. Some sensors might not be available at all times, while others should only be used sparingly to preserve power. In our experience, it is often advisable to temporarily reduce the influence of slower sensors such as WiFi while the user is walking, in order not to disrupt the more accurate position tracking via PDR. Taking knowledge about the context into account, it is possible to refrain from potentially expensive calculations on raw sensor data. Instead, the indoor positioning algorithm has to prioritise the recognition of user activities that given the context and the navigation instruction are most plausible to be performed at any point of time. Lastly, in case no accurate position estimate is computable based on device sensors, the system could enter into a dialog with the user, e.g. asking whether a landmark expected from the context is actually nearby.

As far as the practical purpose of implementing pedestrian navigation systems is concerned we can state that our experiments indicate that integrating context helps to eliminate the limits of indoor positioning exclusively based on sensor data. In particular, integrating metric sensor data and context information with landmarks and navigation instructions based on them help to model the cooperation between users and their navigation system and lead to cognitively adequate and easy-to-use navigation strategies.

In our future work, we will also work on improving SHS and for this purpose will evaluate hybrid system architectures with additional positioning data e.g. from WiFi or BLE in order to analyse their performance to reinitialise PDR after a complete failure. Finally, the observed variations in the magnetic field that heavily influence the compass of smartphones can perhaps be turned into an advantage: we will investigate whether locations exhibit characteristic patterns of magnetic field data that can be exploited for distinguishing locations among each other. Furthermore, we will develop an real-time approach to predict critical PDR errors online and investigate whether in this way we can automatically switch to a classifier for the expected activity in order to further reduce area match failures.

References

Abdelnasser, Heba, Reham Mohamed, Ahmed Elgohary, Moustafa Farid Alzantot, He Wang, Souvik Sen, Romit Roy Choudhury, and Moustafa Youssef. 2016. "SemanticSLAM: Using Environment Landmarks for Unsupervised Indoor Localization." *IEEE Transactions on Mobile Computing* 15 (7): 1770–1782.

Basso, S., G. Frigo, and G. Giorgi. 2015. "A smartphone-based indoor localization system for visually impaired people." In *2015 IEEE International Symposium on Medical Measurements and Applications (MeMeA) Proceedings*, May, 543–548.

Beauregard, Stéphane, Martin Klepal, et al. 2008. "Indoor PDR performance enhancement using minimal map information and particle filters." In *Position, Location and Navigation Symposium, 2008 IEEE/ION*, 141–147. IEEE.

Bohannon, Richard W., and A. Williams Andrews. 2011. "Normal walking speed: A descriptive meta-analysis." *Physiotherapy* 97 (3): 182–189.

Davidson, Pavel, and Robert Piché. 2016. "A Survey of Selected Indoor Positioning Methods for Smartphones." *IEEE Communications Surveys & Tutorials* .

De La Osa, Carlos Martínez, Grigorios G. Anagnostopoulos, Mauricio Togneri, Michel Deriaz, and Dimitri Konstantas. 2016. "Positioning evaluation and ground truth definition for real life use cases." *2016 International Conference on Indoor Positioning and Indoor Navigation, IPIN 2016* .

Deng, Zhihong, Yun Cao, Pengyu Wang, and Bo Wang. 2018. "An Improved Heuristic Drift Elimination Method for Indoor Pedestrian Positioning." *Sensors* 18 (6). http://www.mdpi.com/1424-8220/18/6/1874.

Dez, L. E., A. Bahillo, S. Bataineh, A. D. Masegosa, and A. Perallos. 2016. "Enhancing improved heuristic drift elimination for step-and-heading based pedestrian dead-reckoning systems." In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Aug, 4415–4418.

Ebner, Frank, Toni Fetzer, Frank Deinzer, Lukas Köping, and Marcin Grzegorzek. 2015. "Multi sensor 3D indoor localisation." *2015 International Conference on Indoor Positioning and Indoor Navigation (IPIN)* .

Google. 2017. "Sensor types." https://source.android.com/devices/sensors/sensor-types, Accessed: 2018-09-03.

Guo, H., M. Uradzinski, H. Yin, and M. Yu. 2015. "Indoor positioning based on foot-mounted IMU." *Bulletin of the Polish Academy of Sciences Technical Sciences* 63: 629–634. 3.

Harle, Robert. 2013. "A Survey of Indoor Inertial Positioning Systems for Pedestrians." *IEEE Communications Surveys & Tutorials* 15 (3): 1281–1293.

Hase, K., and R. B. Stein. 1999. "Turning Strategies During Human Walking." *Journal of Neurophysiology* 81 (6): 2914–2922.

Herrera, J. C. A., P. G. Plger, A. Hinkenjann, J. Maiero, M. Flores, and A. Ramos. 2014. ""Pedestrian indoor positioning using smartphone multi-sensing, radio beacons, user positions probability map and Indoor-OSM floor plan representation"." In *2014 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Oct, 636–645.

Hilsenbeck, Sebastian, Dmytro Bobkov, Georg Schroth, Robert Huitl, and Eckehard Steinbach. 2014. "Graph-based Data Fusion of Pedometer and WiFi Measurements for Mobile Indoor Positioning." *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* .

Hoffmann, Eike Jens, Lorenz Schauer, Mirco Schönfeld, and Maximilian Kraus. 2017. "Robust pedestrian dead reckoning using anchor point recalibration." In *2017 International Conference on Indoor Positioning and Indoor Navigation, IPIN 2017*, Sept, 1–7.

Kattenbeck, Markus. 2016. "Empirically Measuring Salience of Objects for Use in Pedestrian Navigation." Dissertation, University Regensburg.

Klepal, Martin, Stéphane Beauregard, et al. 2008. "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*, 79–84. ACM.

Koivisto, Mike, Henri Nurminen, Simo Ali-Loytty, and Robert Piche. 2015. "Graph-based map matching for indoor positioning." In *Information, Communications and Signal Processing (ICICS), 2015 10th International Conference on*, 1–5. IEEE.

Liao, Lin, D. Fox, J. Hightower, H. Kautz, and D. Schulz. 2003. "Voronoi tracking: location estimation using sparse and noisy sensor data." In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, Vol. 1, Oct, 723–728.

Link, Jo Agila Bitsch, Paul Smith, Nicolai Viol, and Klaus Wehrle. 2013. "Accurate Map-based Indoor Navigation on the Mobile." *J. Locat. Based Serv.* 7 (1): 23–43.

Lymberopoulos, Dimitrios, and Jie Liu. 2017. "The Microsoft Indoor Localization Competition: Experiences and Lessons Learned." *IEEE Signal Processing Magazine* 34 (5): 125–140.

Lymberopoulos, Dimitrios, Jie Liu, Xue Yang, Romit Roy Choudhury, Vlado Handziski, and Souvik Sen. 2015. "A Realistic Evaluation and Comparison of Indoor Location Technologies: Experiences and Lessons Learned." In *Proceedings of the 14th international conference on information processing in sensor networks*, 178–189. ACM.

Magdy, N., M. A. Sakr, T. Mostafa, and K. El-Bahnasy. 2015. "Review on trajectory similarity measures." In *2015 IEEE Seventh International Conference on Intelligent Computing and*

Information Systems (ICICIS), Dec, 613–619.

Müller, Manuel, Christina Ohm, Florin Schwappach, and Bernd Ludwig. 2017. "The path of least resistance." *KI - Künstliche Intelligenz* 31 (2): 125–134.

Muro-de-la Herran, Alvaro, Begonya Garcia-Zapirain, and Amaia Mendez-Zorrilla. 2014. "Gait Analysis Methods: An Overview of Wearable and Non-Wearable Systems, Highlighting Clinical Applications." *Sensors (Basel, Switzerland)* 14 (2): 3362–3394.

Nandikolla, Vidya K, Robin Bochen, Steven Meza, and Allan Garcia. 2017. "Experimental Gait Analysis to Study Stress Distribution of the Human Foot." *Journal of Medical Engineering* 2017: 3432074. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC5694576/.

Novak, Domen, Maja Goršič, Janez Podobnik, and Marko Munih. 2014. "Toward Real-Time Automated Detection of Turns during Gait Using Wearable Inertial Measurement Units." *Sensors (Basel, Switzerland)* 14 (10): 18800–18822. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4239865/.

Oberg, T, A Karsznia, and K Oberg. 1993. "Basic gait parameters: Reference data for normal subjects, 10-79 years of age." *Journal of rehabilitation research and development* 30 (2): 210–23.

Ohm, Christina, Manuel Müller, and Bernd Ludwig. 2015. "Displaying landmarks and the user's surroundings in indoor pedestrian navigation systems." *Journal of Ambient Intelligence and Smart Environments* 7 (5): 635–657.

Pham, Duy Duong, and Young Soo Suh. 2016. "Pedestrian Navigation Using Foot-Mounted Inertial Sensor and LIDAR." *Sensors* 16 (1).

Potortì, Francesco, Paolo Barsocchi, Michele Girolami, Joaquín Torres-Sospedra, and Raúl Montoliu. 2015. "Evaluating Indoor Localization Solutions in Large Environments Through Competitive Benchmarking: The EvAAL-ETRI Competition." In *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*, 1–10. IEEE.

Pulkkinen, Teemu, and Johannes Verwijnen. 2015. "Evaluating indoor positioning errors." In *Information and Communication Technology Convergence (ICTC), 2015 International Conference on*, 167–169. IEEE.

Romanovas, Michailas, Vadim Goridko, Lasse Klingbeil, Mohamed Bourouah, Ahmed Al-Jawad, Martin Traechtler, and Yiannos Manoli. 2013. *Pedestrian Indoor Localization Using Foot Mounted Inertial Sensors in Combination with a Magnetometer, a Barometer and RFID*, 151–172. Berlin, Heidelberg: Springer Berlin Heidelberg.

Rossi, Mirco, Julia Seiter, Oliver Amft, Seraina Buchmeier, and Gerhard Tröster. 2013. "Room-Sense: An Indoor Positioning System for Smartphones Using Active Sound Probing." In *Proceedings of the 4th Augmented Human International Conference*, AH '13, New York, NY, USA, 89–95. ACM. http://doi.acm.org/10.1145/2459236.2459252.

Sprager, Sebastijan, and Matjaz B. Juric. 2015. "Inertial Sensor-Based Gait Recognition: A Review." *Sensors* 15 (9): 22089–22127.

Sreenivasa, Manish N., Ilja Frissen, Jan L. Souman, and Marc O. Ernst. 2008. "Walking along curved paths of different angles: the relationship between head and trunk turning." *Experimental Brain Research* 191 (3): 313–320. https://doi.org/10.1007/s00221-008-1525-3.

Susi, Melania, Valérie Renaudin, and Gérard Lachapelle. 2013. "Motion Mode Recognition and Step Detection Algorithms for Mobile Phone Users." *Sensors (Basel, Switzerland)* 13 (2): 1539–1562.

Thrun, Sebastian, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press.

Torres-Sospedra, Joaquín, Adriano Moreira, Stefan Knauth, Rafael Berkvens, Raul Montoliu, Oscar Belmonte, Sergio Trilles, et al. 2017a. "A Realistic Evaluation of Indoor Positioning Systems Based on Wi-Fi Fingerprinting: The 2015 EvAAL–ETRI Competition." *Journal of Ambient Intelligence and Smart Environments* 9 (2): 263–279.

Torres-Sospedra, Joaqun, Antonio R. Jimnez, Stefan Knauth, Adriano Moreira, Yair Beer, Toni Fetzer, Viet-Cuong Ta, et al. 2017b. "The Smartphone-Based Offline Indoor Location Competition at IPIN 2016: Analysis and Future Work." *Sensors* 17 (3). http://www.mdpi.com/1424-8220/17/3/557.

Torres-Sospedra, Joaqun, Antonio R. Jimnez, Adriano Moreira, Toms Lungenstrass, Wei-Chung Lu, Stefan Knauth, Germn Martn Mendoza-Silva, et al. 2018. "Off-Line Evaluation of Mobile-Centric Indoor Positioning Systems: The Experiences from the 2017 IPIN Competition." *Sensors* 18 (2). http://www.mdpi.com/1424-8220/18/2/487.

Verma, S., R. Omanwar, V. Sreejith, and G. S. Meera. 2016. "A smartphone based indoor navigation system." In *2016 28th International Conference on Microelectronics (ICM)*, Dec, 345–348.

Waqar, Wasiq, Yuanzhu Chen, and Andrew Vardy. 2016. "Smartphone positioning in sparse Wi-Fi environments." *Computer Communications* 73: 108 – 117.