A Comparison of Similarity Measures for Localization with Passive RFID Fingerprints

Philipp Vorst, Andreas Zell
Computer Science Department, University of Tübingen, Germany

Abstract

We present a novel approach to self-localization with passive RFID fingerprints using vector space similarity measures and weighted k-nearest neighbors (WKNN). Embedded in a particle filter, our technique provides robust and accurate position estimates. This is shown through various experiments with a mobile robot. The formulation of the observation model guarantees applicability to any type of off-the-shelf RFID reader which counts RFID tag detections. Different similarity measures are investigated, and a new measure is proposed.

1 Introduction

Due to its high-fidelity data association, radio-frequency identification (RFID) has become a valuable sensor in robotics. Several studies have shown how to localize robots – and mobile devices in general – at an accuracy far better than just the cell of origin defined by a detected transponder. The developed techniques have made it possible to cope with the inherent noisiness of long-range passive RFID measurements, and they have tackled the problem that RFID tags are merely identified, without information about bearing and range to them.

In this paper we revisit self-localization with long-range (UHF) passive RFID. We pursue a combination of particle filtering for position estimation and fingerprinting for modeling observations. As usual, the term fingerprinting means that raw observations (RFID measurements) are assumed to explicitly characterize the locations where they were recorded. An explicit model of the interaction between sensor and environment is not required. The derived observation model has the advantage that it is simple and applicable to all long-range RFID readers of the standard EPC Class 1 Gen. 2. This standard is currently used for carton and pallet tagging in industry and of growing interest for item-level tagging. Our approach compares RFID measurements to reference fingerprints for inferring pose estimates, using vector similarity measures. We examine different measures for the core fingerprint retrieval and assessment task. Some measures are well-known and broadly used, others less common or entirely novel.

This paper is organized as follows: In Sect. 2 we review related work. Thereafter, in Sect. 3, we briefly describe the technology and the consequences for the localization task. Our proposed approach is presented in Sect. 4. In Sect. 5 we describe and briefly discuss vector space similarity measures. Experiments with a mobile robot are evaluated in Sect. 6, before we draw conclusions in Sect. 7.

2 Related Work

Localization has been an active field of research for several years. Among countless pieces of scientific work, we here focus on three types: RFID-based localization, positioning with fingerprints using wireless identification, and fingerprinting-based localization with mobile robots. With regard to the latter category, the *visual appearance* of a scene is probably the most widely used type of fingerprint: A robot is localized using visual features extracted from camera images which are compared to reference views. In contrast to stereo vision or structure from motion, the coordinates of landmarks corresponding to recognized features need not be reconstructed. Examples are [17, 21], in which local integral invariants are employed, in [21] combined with Monte Carlo localization. Ulrich and Nourbakhsh also compared similarity measures for appearance-based place recognition using color histograms [18]. Recent examples of the more general approach without prior environment map, known as appearance-based simultaneous localization and mapping (SLAM), are the works by Cummins and Newman [4], Konolige and Agrawal [9], or the minimalistic approach by Andreasson et al. [1]. All these works more or less build on techniques developed for the more general field of content-based image retrieval (CBIR). A comparison of similarity measures for CBIR is presented in [13], for instance.

Wireless fingerprinting approaches utilize the identification mechanisms of wireless standards (WiFi, Bluetooth, GSM etc.) which can be regarded as *active* (=battery-powered) RFID. Fingerprints commonly consist of the addresses of detected devices and information about signal strength or link quality. Sample approaches using signal strength (RSSI) are Radar [3] and its extension to Markov localization by Ladd et al. [10]. Li et al. [11] compared WKNN and Markov localization based on WLAN signal

strength measurements. Examples of fingerprinting with other proprietary active RFID systems are [14, 22].

With regard to fingerprinting-based localization with RFID, Lim and Zhang developed a deterministic approach for passive RFID tags on the ceiling [12]. Joho et al. localized a shopping cart with a signal strength map which was recorded with a passive RFID reader [8]. In [15], we presented a probabilistic fingerprinting approach based on local regression and Bayes estimation of detection rates, embedded in a particle filter. This approach requires knowledge about how often the RFID reader issues inquiries.

3 Long-range Passive RFID

Radio-frequency identification (RFID) is a technology originally designed for the contactless identification of RFID-tagged objects via radio waves. An RFID reader emits an electromagnetic field which enables passive (=batteryless) transponders in read range to send back their unique identifier. According to the industry standard EPC C1G2, transponder IDs consist at least of a unique 96-bit number, the electronic product code (EPC). Long-range systems in the UHF band have a read range of several meters (sometimes up to $10\,\mathrm{m}$) at full power level ($2\,\mathrm{W}$ EIRP). Since tag IDs are unique and the reader detects transmission errors, it is virtually impossible to recognize the wrong object. This property is valuable for robotics: RFID tags can be used as navigation landmarks with perfect data association.
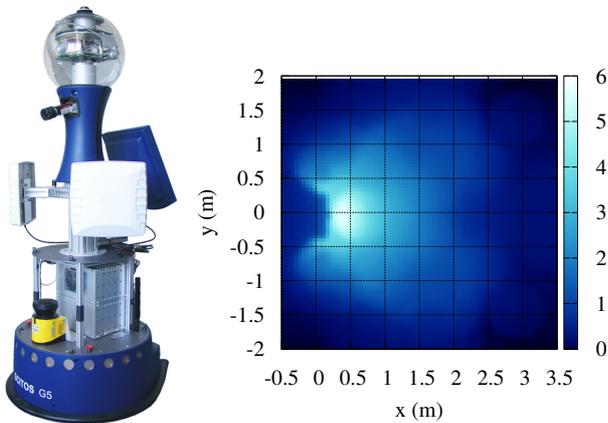
Figure 1: Left: The experimental platform, a MetraLabs SCITOS G5, with an onboard passive UHF RFID reader (reader antennas are white) and a laser range finder (yellow) for ground truth. Right: A model of detection rates for the employed RFID system. This model is *not* required for our fingerprinting approach and just shown to depict detection characteristics. Right in front of the reader antenna a single transponder is identified six times per measurement on average, and only once in a distance of approx. 2.5 m.

Passive RFID also has some shortcomings, which we address in the work at hand: First, detection rates are noisy,

and transponders within theoretical read range are frequently not detected. Moreover, standard-compliant RFID systems can only determine the presence of a tag; the distance or bearing to the tag are not disclosed. This is in opposition to active RFID, where signal strength indicates the distance to a tag. Detection rates, however, do depend on the relative position of the RFID tag with respect to the antenna of the RFID reader. A model of this relationship is visualized in **Fig. 1**. Such a sensor model can be utilized for localization purposes, as shown by several researchers [7, 8] if, besides the model, the positions of static RFID tags in the environment are known. Other factors such as materials (metal or water) in the vicinity of RFID reader and tag reveal major influence on the possibility to successfully read the tag ID, too, and can hardly be modeled.

For this reason, we again pursue a fingerprinting approach to localization in this paper: Instead of using a model of the detection rates of the reader, we infer the pose of the robot directly from reference RFID measurements recorded in a prior mapping phase. The idea is to make no assumptions about how and where RFID transponders are distributed. Aiming at future application scenarios, tags will most likely mark shelves and products to the sides of corridors in stores, and our localization technique utilizes the given RFID infrastructure. The learnt distribution of measurements also implicitly contains the hardly modeleable environment characteristics, which promises better localization accuracy. During a calibration/mapping stage, the robot records RFID detections at several positions. A laser range finder provides reference positions in our case. Later, in the localization stage, the pose of the robot is estimated by comparing the current RFID measurement to the fingerprints from the calibration step.

Passive RFID systems always report tag detection counts, i.e., for each tag which was detected, the reader reports how often. In autonomous mode, the RFID reader repeatedly perform inquiries in the background. Only after a (usually fixed) time interval T, the list of detection counts is reported. The actual number of read attempts made in T depends on several factors, e.g., the number of transponders close to the RFID reader, and may be unknown. This information, however, is required in other probabilistic RFID fingerprint approaches [15].

In this paper, we generalize the interpretation of an RFID measurement. Mathematically, we treat RFID tag detection counts as vectors $\mathbf{f}_a = (f_1, \ldots, f_L)$, where $f_i \in \mathbb{N}_0$ counts how often tag i has recently been detected by antenna a. L is the number of tags in the environment, but needs not be known in advance. It typically increases while the robot is moving through the environment. So, technically, we store $\mathbf{f}_a$ as a list and set, when required, $f_l = 0$ for every not observed transponder l. Moreover, contrary to related approaches [12, 15], f_i does not have a predefined upper bound. $\mathbf{f} = (\mathbf{f}_1, \ldots, \mathbf{f}_A)$ finally makes the tuple comprised of tag lists $\mathbf{f}_a$ for all A antennas.

4 Approach

4.1 Mapping/Calibration

Every fingerprinting approach starts with a mapping stage during which fingerprints are collected at known reference positions. As common, our robot (see **Fig. 1**) possesses a laser range finder: Hence, it can be steered through the environment while recording RFID measurements and annotating them with the latest pose estimates obtained via classical laser-based localization [6]. The result of the mapping phase is a set $\mathbf{m} = \{\mathbf{F}_1, \mathbf{F}_2, \ldots\}$ of reference fingerprints $\mathbf{F}_i = (\mathbf{f}_i, \mathbf{x}_i)$, containing RFID measurement $\mathbf{f}_i$ and the corresponding pose $\mathbf{x}_i = (x_i, y_i, \theta_i)$ of the robot. (x_i, y_i) are the coordinates of the robot in a global frame of reference, and θ_i is the global heading of the robot.

The mapping stage can be tedious if intensive human intervention is required. If transponder positions are not known, however, model-based approaches require thorough mapping as well. Moreover, reference positions from another positioning system are required. Solutions to these issues are provided by exploration and SLAM techniques (e.g., [20]), but are beyond the scope of this work.

4.2 Localization

4.2.1 Particle Filtering

For localizing the robot, we employ a particle filter [2] (also known as Monte Carlo localization [6]): At each time step t the pose of the robot $\mathbf{x}_t$ is modelled by a probability density function (pdf) over the space of locations. This pdf is approximated by a set of n samples (particles). Each sample consists of a pose hypothesis $\mathbf{x}_t^i = (x_t^i, y_t^i, \theta_t^i)$ (2D position+heading) and an importance weight w_t^i. The initial distribution, $p(\mathbf{x}_0)$, is usually assumed uniform over the state space. The filtering algorithm then iteratively performs three steps:

1. *Prediction*: Given a new odometry reading $\mathbf{o}_t$ since iteration $t-1$, the robot pose at time t is predicted by propagating all particle positions according to a motion model $p(\mathbf{x}_t|\mathbf{o}_t, \mathbf{x}_{t-1})$.

2. *Correction*: A new measurement $\mathbf{g}_t$ leads to a correction of particle weights according to the likelihood function $p(\mathbf{g}|\mathbf{x}, \mathbf{m})$, the observation model:

$$w_t^i = \eta_t \, w_{t-1}^i \, p(\mathbf{g}_t|\mathbf{x}_t^i, \mathbf{m}) \qquad (1)$$

   $\mathbf{m}$ is an environment representation (in our case the set of reference fingerprints from Sect. 4.1). The normalizer η_t ensures that $\sum_{i=1}^n w_t^i = 1$.

3. *Resampling*: A new set of n particles with equal weights $1/n$ is obtained by drawing n times a sample from the old set of particles. The probability of choosing particle i corresponds to its weight w_t^i. An

option is to resample not always, but only if the estimate $\hat{n}_{eff} \approx 1/\left(\sum_{i=1}^n (w_t^i)^2\right)$ of the effective sample size falls below some threshold, e.g. $n/2$. There are several resampling methods (for a comparison we refer to [5]), among which we chose *residual resampling*.

The current pose of the robot can be estimated by $\hat{\mathbf{x}}_t = \sum_{i=1}^n w_t^i \mathbf{x}_t^i$. In this way, particle filtering enables a robot both to localize itself globally and to track its position over time. A major factor of robustness and versatility is that particle filters can deal with arbitrary probability distributions of noise with respect to motions and sensor data. This makes them a good choice when using passive RFID.

4.2.2 Modeling Observations

The observation model is represented by the likelihood $p(\mathbf{g}|\mathbf{x}^i, \mathbf{m})$ for a sample placed at position $\mathbf{x}^i$, given the current measurement $\mathbf{g}$. Analogously to the reference fingerprints, $\mathbf{g} = (\mathbf{g}_1, \ldots, \mathbf{g}_A)$ consists of tag counts for all A antennas. In the following, we derive a likelihood function, based on a k-nearest neighbor search for reference fingerprints and subsequent weighting.

First, all reference fingerprints $\mathbf{f}_i$ are searched which contain at least one tag also detected in $\mathbf{g}$. An inverted index of reference fingerprints is used for this purpose, that is, a mapping from transponder identifiers l to indices of reference fingerprints containing l. Then the similarities $\mathrm{sim}_a(\mathbf{g}_a, \mathbf{f}_{i,a})$ are computed individually for all antennas a, $a = 1, \ldots, A$, by means of similarity measures depicted in Sect. 5. Generally, there are only few basic constraints on the choice of the function sim: First, it should yield nonnegative values. Second, values close to 0 indicate dissimilarity, and the larger the computed value, the larger is the similarity of the two compared RFID measurements.

In the next step we calculate the similarity of the current measurement $\mathbf{g}$ and all reference fingerprints $\mathbf{f}_i$. The similarities at all antennas are integrated by a weighted average:

$$\mathrm{sim}(\mathbf{g}, \mathbf{f}_i) = \sum_{a=1}^A \mathrm{sim}_a(\mathbf{g}_a, \mathbf{f}_{i,a}) \cdot \frac{n(\mathbf{g}_a, \mathbf{f}_{i,a})}{\sum_{a=1}^A n(\mathbf{g}_a, \mathbf{f}_{i,a})} \quad (2)$$

where

$$n(\mathbf{g}_a, \mathbf{f}_{i,a}) := \max(|\mathbf{g}_a|, |\mathbf{f}_{i,a}|)$$

be the maximal number of detected tags in the vectors $\mathbf{g}_a$ and $\mathbf{f}_{i,a}$.

Using sim, we search for the k most similar reference fingerprints $\mathbf{F}_{i_1}, \ldots, \mathbf{F}_{i_k}$.[1]

Then, we compute the likelihood $p(\mathbf{g}|\mathbf{x}, \mathbf{m})$ of observing $\mathbf{g}$ from the pose $\mathbf{x}$ (the positions of the particles), given M

---

[1]If there are only $k' < k$ fingerprints of nonzero similarity, one proceeds with only k' fingerprints.

reference fingerprints in $\mathbf{m}$:

$$p(\mathbf{g}|\mathbf{x}, \mathbf{m}) \tag{3}$$

$$= \sum_{j=1}^{M} p(\mathbf{g}|\mathbf{x}, \mathbf{F}_j) p(\mathbf{F}_j|\mathbf{x}) \tag{4}$$

$$= \sum_{j=1}^{M} \nu_s \mathrm{sim}(\mathbf{g}, \mathbf{f}_j)\, \nu_d \exp\left(-\frac{1}{2} d(\mathbf{x}_j, \mathbf{x})\right) \tag{5}$$

Eq. 4 follows from the law of total probability. In Eq. 5, we model $p(\mathbf{g}|\mathbf{x}, \mathbf{F}_j)$ (with $\mathbf{F}_j = (\mathbf{f}_j, \mathbf{x}_j)$) by the similarity of $\mathbf{g}$ and $\mathbf{f}_j$, normalized with a suitable ν_s. Moreover, $p(\mathbf{F}_j|\mathbf{x})$ is represented by a density depending on the distance between $\mathbf{x}$ and the jth reference fingerprint $\mathbf{f}_j$, again normalized with a suitable ν_d. $d(\cdot)$ is a squared distance assessing both translational and rotational displacement:

$$d(\mathbf{x}, \mathbf{x}_{i_j}) = \frac{(x - x_{i_j})^2}{\sigma_d^2} + \frac{(y - y_{i_j})^2}{\sigma_d^2} + \frac{(\theta \ominus \theta_{i_j})^2}{\sigma_r^2} \tag{6}$$

The $\ominus$ denotes the difference of angles, restricted to the interval $[-\pi, \pi]$. σ_d and σ_r are bandwidth parameters for the translational and the rotational distance components, respectively. The final approximation

$$p(\mathbf{g}|\mathbf{x}, \mathbf{m}) \approx \nu \sum_{j=1}^{k} \mathrm{sim}(\mathbf{g}, \mathbf{f}_{i_j})\, \exp\left(-\frac{1}{2} d(\mathbf{x}_{i_j}, \mathbf{x})\right) \tag{7}$$

builds on the assumption that the k most similar measurements capture most of the likelihood in Eq. 5. We further set $\nu = \nu_s \nu_d$. Note that ν needs not be computed explicitly, because sample weights are normalized after applying the observation model.

5 Similarity Measures

In this section we investigate similarity measures for comparing two RFID measurement vectors (fingerprints) and for weighting particles as described in the previous section. In addition to well-known measures, we examine new similarity measures which take characteristics of passive long-range RFID into account.

5.1 General Considerations

The choice of an adequate similarity measure can have major impact on the localization result. In the technical literature, a variety of similarity measures has been employed for different applications. This already shows that a suitable measure is task-dependant. A similarity measure for long-range passive RFID should take into consideration:

1. *Data association* is known. Hence measurement vectors can be compared in a component-wise fashion. *Cross-bin* similarity measures (such as earth mover's distance or cross-correlation) need not be used.

2. *False-negative detections* frequently occur and must not be overrated. RFID tags in read range may not be detected even if detected by a prior inquiry in the same position.

3. The *number of tag IDs in common* in two compared measurements may be more important than how often each of the tags was counted: Detection rates are typically noisy, but if there are several RFID tags spread over the environment, the overlap of tag identifiers helps to refine the position of the RFID antenna.

5.2 A Selection of Similarity Measures

In **Tab. 1** we list the employed measures, which we detail in the following. All functions are nonnegative and symmetric. That is, $\mathrm{sim}(\mathbf{g}, \mathbf{f}) = \mathrm{sim}(\mathbf{f}, \mathbf{g}) \geq 0$. We need not bound $\mathrm{sim}(\mathbf{f}, \mathbf{g})$, because our approach is scale-independent w.r.t. the similarity space.[2]

5.2.1 Measures of Similarity

Histogram intersection (abbreviated by HIST) is a widespread measure of similarity. It was originally developed to compare color histograms [16] for image retrieval. Given two RFID measurement vectors, it indicates how many tag counts two measurements have in common in each component. This makes the measure partly robust to outliers in single components of measurements

Another classical similarity measure under investigation is the *cosine similarity* (COS). Visually, it represents the cosine of the angle spanned by two vectors. We have already successfully used it for determining loop closure in trajectory estimation [20]. As a variant, we also examine C*H, a function which simply multiplies the values of COS and HIST, with the idea to yield a more distinctive measure combining the advantages of COS and HIST.

The *Bhattacharyya coefficient* (BHA) is similar to the vector dot product, but possesses an additional inner square root. This makes the coefficient less sensitive to larger tag detection counts in single components.

Additionally, we examine a novel measure, which we called *overlap score* (OSC). The idea was to use the cosine similarity, COS, and weight it with the number of tags in common, NCT. The enclosing logarithm introduces non-linearity, with the idea that another tag in common only marginally increases the likeliness that both RFID measurements stem from the same position.

---

[2] The k-nearest neighbors search is scale-independent, and particle weights will be normalized due to Eq. (1).

Similarity measure	Abbreviation	Formula	Range
Cosine similarity	COS	$\text{sim}_{\text{COS}}(\mathbf{f}, \mathbf{g}) = \dfrac{\sum_{l=1}^{L} f_l g_l}{\sqrt{\sum_{l=1}^{L}(f_l)^2} \cdot \sqrt{\sum_{l=1}^{L}(g_l)^2}}$	$[0, 1]$
Histogram intersection	HIST	$\text{sim}_{\text{HIST}}(\mathbf{f}, \mathbf{g}) = \sum_{l=1}^{L} \min(f_l, g_l)$	$[0, \infty)$
Bhattacharyya coefficient	BHA	$\text{sim}_{\text{BHA}}(\mathbf{f}, \mathbf{g}) = \sum_{l=1}^{L} \sqrt{f_l g_l}$	$[0, \infty)$
Overlap score	OSC	$\text{sim}_{\text{OSC}}(\mathbf{f}, \mathbf{g}) = \log(1 + \text{sim}_{\text{NCT}}(\mathbf{f}, \mathbf{g})\,\text{sim}_{\text{COS}}(\mathbf{f}, \mathbf{g}))$	$[0, \infty)$
Minkowski distance	L_p	$d_p(\mathbf{f}, \mathbf{g}) = \left(\sum_l |f_l - g_l|^p\right)^{\frac{1}{p}}$	$[0, \infty)$
Hellinger distance	HD	$d_{\text{HD}}(\mathbf{f}, \mathbf{g}) = \sqrt{\sum_{l=1}^{L}(\sqrt{f_l} - \sqrt{g_l})^2}$	$[0, \infty)$
χ^2 statistics	CHI	$d_{\text{CHI}}(\mathbf{f}, \mathbf{g}) = \sum_{l=1}^{L}(f_l - \mu_l)^2/\mu_l, \quad \mu_l = \frac{f_l + g_l}{2}$	$[0, \infty)$
Jeffrey divergence	JD	$d_{\text{JD}}(\mathbf{f}, \mathbf{g}) = \sum_{l=1}^{L}(f_l \log(f_l/\mu_l) + g_l \log(g_l/\mu_l))$	$[0, \infty)$
Dot (or scalar) product	DOT	$\text{sim}_{\text{DOT}}(\mathbf{f}, \mathbf{g}) = \sum_{l=1}^{L} f_l g_l$	$[0, \infty)$
Number of common tags	NCT	$\text{sim}_{\text{NCT}}(\mathbf{f}, \mathbf{g}) = |\{l \mid f_l g_l > 0, l = 1, \ldots, L\}|$	$[0, \infty)$

Table 1: Investigated similarity measures sim for comparing two tag lists $\mathbf{f}$ and $\mathbf{g}$.

5.2.2 Measures of Dissimilarity

Vector distance measures assess only the *dissimilarity* of vectors. Still, they can easily be transformed into similarity functions: If $d \in [0, \infty)$ is a distance value, then

$$s(d) = \frac{1}{d + \varepsilon} \in (0, \varepsilon^{-1}], \quad \varepsilon > 0 \qquad (8)$$

is a similarity value. This is not the only method of how to transform distance to similarity, even for the case treated here that d is quasi-unbounded. Another transformation is $s(d) = \exp(-\tau \cdot d)$ for some positive real τ. We found, however, that the resulting behavior is quite sensitive to the values of τ, because the function converges to zero quickly, and committed ourselves to Eq. 8. We should underline that also the constant ε does have impact on both the resulting values of s and the final localization accuracy. Theoretically, the variable should thus be subject to optimization. We think, however, that $\varepsilon = 1$ is an intuitive, well justifiable choice, since $s(0) = 1$.

A widespread class of distances are the *Minkowski distances* (or L_p-*norm*). *Euclidean distance* ($p = 2$), L_2, and *Manhattan/city block distance* ($p = 1$), L_1, are two special cases which are frequently used in the fingerprinting literature. However, when it comes to comparing RFID measurements, one can find undesirable mathematical properties. As one example, let us consider three measurements $\mathbf{f}$, $\mathbf{g}$, and $\mathbf{h}$, where the counts of the jth tag are $f_j = 0$, $g_j = 1$, $h_j = 2$, and for simplicity let all other tag counts equal some value c. Then $d_p(\mathbf{f}, \mathbf{g}) = d_p(\mathbf{g}, \mathbf{h})$. This is counterintuitive, because given a measurement $g_j = 1$, the tag count $h_j = 2$ is a stronger indication that $\mathbf{h}$ was recorded in the same area as $\mathbf{g}$ than that $\mathbf{f}$ was recorded in the same place, because in $\mathbf{f}$ tag j was not detected.

An alternative against this background is the *Hellinger distance* (HD, the square root of the *squared chord* measure [13]). It differs from the Euclidean distance by a non-linearity in the vector components, which solves the above paradox.

Another option is to scale deviations with the inverse of the mean of two measurements. This is captured by the χ^2 *statistics* (CHI)[3]. In probabilistic terms, d_{CHI} represents the likeliness of $\mathbf{f}$ being drawn from a distribution $\mathbf{g}$.

The fourth distance measure, also known from information theory, is the *Jeffrey divergence* (JD). It represents the symmetric, numerically stable variant of the Kullback-Leibler divergence. We conclude this section with the remark that χ^2 statistics and JD are known to work well if a large number of observations are available.

5.2.3 Benchmark Measures

In order to investigate the difference in performance between the rather sophisticated, traditional similarity measures above and other rather simple functions, we compare our results to two benchmark functions: The first benchmark measure is the vector *dot product* (or *scalar product*), DOT. In our context the dot product is useful because the more detected tags in common ($f_i, g_i > 0$) and the greater the tag counts in both components f_i and g_i under investigation, the higher the similarity. Moreover, the dot product is appealing because of its simplicity of computation. A weakness, on the other hand, is that measurement pairs yielding the same product $f_l g_l$ in a component will receive the same similarity value. For instance, $\text{sim}_{\text{DOT}}((1), (4)) = \text{sim}_{\text{DOT}}((2), (2)) = 4$. Note that the Bhattacharyya coefficient reveals the same issue, but it is less sensitive to larger tag detection counts in single components due to the inner square root. Moreover, it may be undesirable that $\text{sim}_{\text{DOT}}((0), (c)) = 0$, independent of c (a property that holds for all presented measures of similarity, too).

As another non-classical benchmark we utilize the *number of common tags* (NCT) in two measurements. In opposition to the other measures, the number of detections of a specific RFID tag is quantified in a boolean fashion only. Although potentially valuable information is ignored, that

---

[3]In the formula of χ^2 statistics, we set the ith summand zero if $f_i = g_i = 0$.

mechanism makes the measure robust to outliers in the detection values. In [19], we presented a loop closure model for SLAM which was based on NCT. There, it served to recognize previously visited places.

6 Experiments

Setup We conducted experiments with an indoor mobile robot of type SCITOS G5 by MetraLabs (Fig. 1). The robot is equipped with a $270°$ laser range finder for ground truth positioning and an Elatec SR-113 RFID reader. Two circularly polarized UHF antennas are connected to the RFID reader. They scan along horizontal axes spanning an angle of approx. $90°$ in the xy plane.

In the environment depicted in **Fig. 2** (a corridor with adjacent hall) we attached passive RFID tags (type Alien Technology Squiggle, ISO/IEC 18000-6C) to walls, at different heights between the floor and the height of the RFID antennas (0.8 m). The installation was intendedly not overly systematic, besides that we tried to spread tags roughly in a balanced distribution. We tested two different transponder densities. The first density corresponds coarsely to distances of 1.0-2.0 m between each pair of neighboring tags, the second density to an average of 0.5-1.0 m.
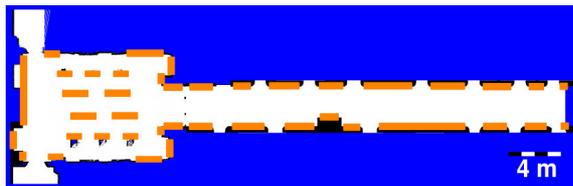
Figure 2: The experimental environment has a traversable space of approx. $160\,\mathrm{m}^2$. Tags were attached to walls in the orange/gray areas.

For each density, we manually steered the robot on 5 different paths, each of which comprises 1000 RFID measurements and corresponding reference positions. The RFID reader was configured to transmit at full power level (1 W EIRP) and to supply the results of RFID inquiries every $T = 0.5\,s$. The mean (and maximum) numbers of detected transponders per inquiry were 8.5 (max. 27) and 4.5 (max. 16) for the higher and the lower tag density, respectively. At the same time, the mean (and maximum) detection counts per detected transponders were 4.1 (max. 24) and 6.6 (max. 26), respectively. In all subsequent experiments, we performed cross-validation by picking the datasets of two recorded paths, random sampling of M reference measurements, and running the particle filter five times on each of the remaining three recorded paths.

Results **Fig. 3** shows the *tracking* results (i.e., the initial pose is known) for the similarity measures from Sect. 5. The four subfigures illustrate the effect of the two different transponder densities and two different numbers of ref-

erence fingerprints ($M = 500, 2000$), while the parameter k was varied. Each outcome is based on 300 single runs of a particle filter with 1000 samples. The parameters $\sigma_{xy} = 0.5$ and $\sigma_\theta = 0.3$ were fixed.

For the higher tag density and $M = 2000$ reference fingerprints, all similarity measures yielded roughly the same accuracy. For $k = 16$, the mean errors were below 0.25 m, with standard deviations of approx. 0.025 m. HIST performed best (0.237 m $\pm$ 0.025 m), but not significantly better. For the case that the tag density is sufficient, also the benchmark measures, DOT and NCT, yield surprisingly good accuracy.

If the transponder density or the number of reference fingerprints is decreased, tracking errors and variances increase, which corresponds to intuition. On average, the tracking accuracy of the measures of dissimilarity (L_1, L_2, HD, CHI, and JD) seems to degrade faster than the accuracy of the measures of similarity (COS, HIST, BHA, OSC, and the product of COS and HIST), which is documented by larger means and variances. In our experiments, especially the L_1 and the L_2 norm were sensitive to tag density and number of reference measurements.

Results for *global localization* (i.e., with unknown initial position) under variation of the number of reference fingerprints are illustrated in **Fig. 4**. In this case, we used 2000 samples to account for the increased uncertainty in the first localization iterations. The mean ad-hoc localization error (that is, the error after evaluating the first observation) was 0.945 m. The error thereafter converged to an average of 0.254 m $\pm$ 0.037 m.
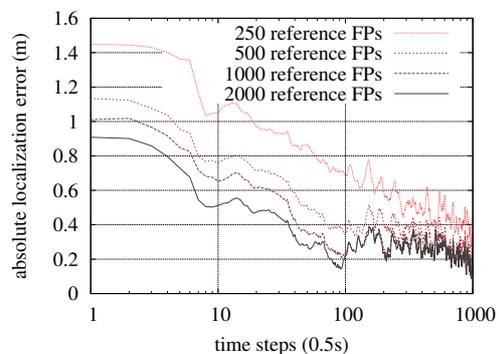
Figure 4: Localization error over time (measure OSC, $k = 16$, 2000 particles), averaged over all experiment runs, for different numbers of reference fingerprints (FPs). The horizontal axis has logarithmic scale in order to emphasize the first localization steps.

The influence of the bandwidth parameters σ_d and σ_r (Eq. 6) on the tracking results is illustrated in **Fig. 5**. For these measurements, we ran a particle filter with 1000 samples, performing cross-validation as above. We employed histogram intersection, 2000 reference fingerprints were used, k was 16, and σ_d and σ_r were varied.

The setting $\sigma_d = 0.5$ and $\sigma_r = 0.3\,(17°)$ yielded the smallest mean error. The differences in the mean errors obtained
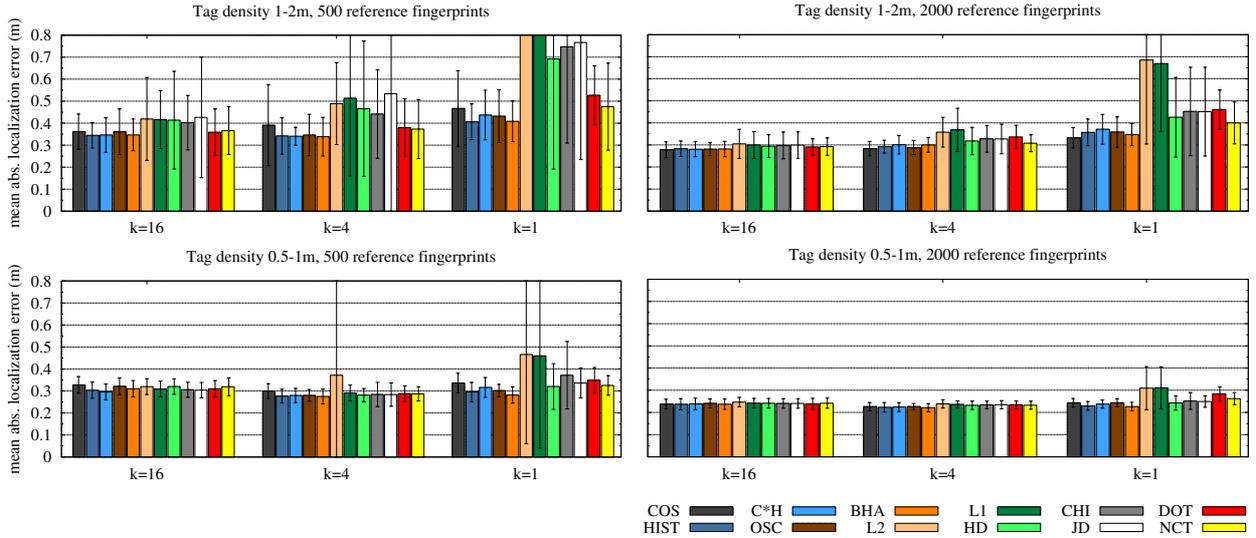
Figure 3: Mean absolute tracking errors for the different similarity measures, tag densities, and choices of k and the number of reference fingerprints (top row lower tag density, bottom row higher tag density; left column 500 reference fingerprints, right column 2000 reference fingerprints). The mean values of L_1 and L_2 in the top left figure are approx. 1.15 m and cut for the sake of clarity.

are, however, not statistically significant. Besides the case $\sigma_d = 1.0$ and $\sigma_r = 0.15$, all means differ by at most 3 cm only. Consequently, even in the case of slightly suboptimally chosen bandwidth parameters, the localization accuracy can be expected to not degrade considerably.
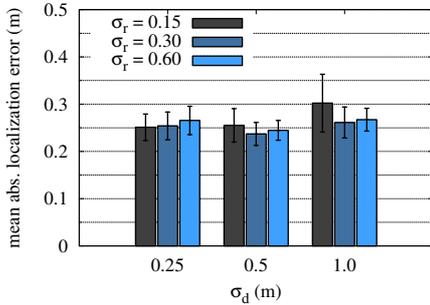
Figure 5: Means and standard deviations of tracking results for varying bandwidths σ_d (horizontal axis) and σ_r (bars).

Discussion The insight that – under the better conditions – the choice of similarity measure seems irrelevant was unexpected, since in many other applications, the choice of measure does matter. The observed invariance to the similarity measure, however, allows the programmer to choose the function depending on other criteria such as computational complexity. On the other hand, in the cases in which fewer information is available – if fewer RFID tags or fewer reference measurements are available –, the performances of the compared similarity measures do differ. Fortunately, if RFID tags are used on a large scale (in supermarkets, for instance), the tag densities will be even

higher than our densest distribution. This promises even better localization accuracy. We did not further increase the transponder density because we think that if a human is required to prepare an environment, the chosen density can quickly be installed.

We conclude this section with the note that all estimations can be performed in real-time. Each correction step took less than 9 ms on a 3 GHz PC for 1000 particles, $k = 16$, and $M = 2000$.

7 Conclusion

The presented approach combines location fingerprinting with the filtering of passive RFID measurements. Advantages are the high accuracy due to location-specific measurements and the universal applicability of the approach. In a number of experiments, we achieved global localization and tracking errors of approx. 0.25 m or even better. These results are comparable to other fingerprinting approaches to RFID-based localization [15]. At the same time, the implementation of the presented novel technique is comparatively simple. We compared different functions and observed that the measures of similarity, as opposed to distance measures, reliably achieved accurate localization. The observation model performs a k-nearest neighbor search in similarity space. Computationally, it scales also for larger environments because of the efficient inverse index for reference fingerprints. The requirement of a reference positioning system and the potentially time-consuming training are drawbacks, as with all fingerprinting approaches. They can be tackled by further automation based on exploration and SLAM algorithms (e.g., [20]).

As part of future work, we plan to investigate the impact of transponder relocations on the localization accuracy.

References

[1] H. Andreasson, T. Duckett, and A. J. Lilienthal. Mini-SLAM: Minimalistic visual SLAM in large-scale environments based on a new interpretation of image similarity. In *Proc. 2007 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 4096–4101, 2007.

[2] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.

[3] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proc. IEEE Infocom*, volume 2, pages 775–784, 2000.

[4] M. Cummins and P. Newman. FAB-MAP: Probabilistic localization and mapping in the space of appearance. *The Int. Journal of Robotics Research*, 27(6):647–665, June 2008.

[5] R. Douc, O. Cappé, and E. Moulines. Comparison of resampling schemes for particle filtering. In *4th Int. Symposium on Image and Signal Processing and Analysis (ISPA)*, Zagreb, Croatia, September 2005.

[6] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo localization: Efficient position estimation for mobile robots. In *Proc. Sixteenth National Conf. on Artificial Intelligence (AAAI)*, 1999.

[7] D. Hähnel, W. Burgard, D. Fox, K. P. Fishkin, and M. Philipose. Mapping and localization with RFID technology. In *Proc. 2004 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 1015–1020, 2004.

[8] D. Joho, C. Plagemann, and W. Burgard. Modeling RFID signal strength and tag detection for localization and mapping. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 3160–3165, Kobe, Japan, 2009.

[9] K. Konolige and M. Agrawal. FrameSLAM: From bundle adjustment to real-time visual mapping. *IEEE Transactions on Robotics*, 24(5):1066–1077, 2008.

[10] A. Ladd, K. Bekris, G. Marceau, A. Rudys, L. Kavraki, and D. Wallach. Robotics-based location sensing using wireless ethernet. Technical Report TR02-393, Department of Computer Science, Rice University, 2002.

[11] B. Li, J. Salter, A. G. Dempster, and C. Rizos. Indoor positioning techniques based on wireless LAN. In *1st IEEE Int. Conf. on Wireless Broadband and Ultra Wideband Communications*, Sydney, Australia, 2006. paper 113, CD-ROM proceedings.

[12] A. Lim and K. Zhang. *Advances in Applied Artificial Intelligence*, volume 4031 of *Lecture Notes in Computer Science*, chapter A Robust RFID-Based Method for Precise Indoor Positioning, pages 1189–1199. Springer, 2006.

[13] H. Liu, D. Song, S. Rüger, R. Hu, and V. Uren. Comparing dissimilarity measures for content-based image retrieval. In *Information Retrieval Technology*, volume 4993/2008 of *Lecture Notes in Computer Science*, pages 44–50, Harbin, China, 2008. Springer.

[14] L. Ni, Y. Liu, Y. C. Lau, and A. Patil. LANDMARC: Indoor location sensing using active RFID. *ACM Wireless Networks (WINET)*, 10(6):701–710, 2004.

[15] S. Schneegans, P. Vorst, and A. Zell. Using RFID snapshots for mobile robot self-localization. In *Proc. 3rd European Conf. on Mobile Robots (ECMR)*, pages 241–246, Freiburg, Germany, 2007.

[16] M. J. Swain and D. H. Ballard. Color indexing. *Int. Journal of Computer Vision*, 7(1):11–32, 1991.

[17] H. Tamimi, A. Halawani, H. Burkhardt, and A. Zell. Appearance-based localization of mobile robots using local integral invariants. In *Proc. 9th Int. Conf. on Intelligent Autonomous Systems (IAS-9)*, pages 181–188, Tokyo, Japan, 2006.

[18] I. Ulrich and I. Nourbakhsh. Appearance-based place recognition for topological localization. In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 1023–1029, San Francisco, USA, 2000.

[19] P. Vorst and A. Zell. Particle filter-based trajectory estimation with passive UHF RFID fingerprints in unknown environments. In *Proc. 2009 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 395–401, St. Louis, USA, 2009.

[20] P. Vorst and A. Zell. Fully autonomous trajectory estimation with long-range passive RFID. In *Proc. 2010 IEEE Int. Conf. on Robotics and Automation (ICRA)*, Anchorage, Alaska, USA, 2010.

[21] J. Wolf, W. Burgard, and H. Burkhardt. Robust vision-based localization by combining an image retrieval system with Monte Carlo localization. *IEEE Transactions on Robotics*, 21(2):208–216, 2005.

[22] K. Yamano, K. Tanaka, M. Hirayama, E. Kondo, Y. Kimuro, and M. Matsumoto. Self-localization of mobile robots with RFID system by using support vector machine. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, volume 4, pages 3756–3761, 2004.