RL-SAR: A Robotic System for Fine-Grained RFID Localization using RL-based Synthetic Aperture Radar

by

Weitung Chen

B.S. Electrical Engineering and Computer Science
Massachusetts Institute of Technology (2021)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2023

Authored by:	Weitung Chen Department of Electrical Engineering and Computer Science July 28, 2023
Certified by:	Fadel Adib Associate Professor, Thesis Supervisor
Accepted by:	Katrina LaCurts Chair, Master of Engineering Thesis Committee

RL-SAR: A Robotic System for Fine-Grained RFID Localization using RL-based Synthetic Aperture Radar

by

Weitung Chen

Submitted to the Department of Electrical Engineering and Computer Science
on July 28, 2023, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Efficient localization of RFID-tagged items is crucial in scenarios that require tracking and managing a large inventory. Current systems for fine-grained RFID localization have shown limitations since they only collect measurements on a pre-defined trajectory or optimize measurement locations for a single tag. Thus, there is a need for an RFID localization system that can autonomously optimize for multiple tags and adaptively relocalize tags with lower confidence to achieve a more precise and efficient localization.

We introduce RL-SAR, an end-to-end autonomous Synthetic Aperture Radar (SAR) based RFID localization system, utilizing a Reinforcement Learning (RL) algorithm to determine the most optimal trajectory for localizing multiple tags. We implemented this system with an antenna moving on a ceiling-mounted 2D track. The core of the system is a RL-based trajectory optimization algorithm for collecting RF measurements. Based on these RF measurements, we developed a data processing pipeline to compute the estimated tag locations along with their confidence metrics, derived from the RF SAR hologram. The RL algorithm leverages confidence metrics associated with the tags and is capable of learning a strategy that minimizes the antenna's traveled distance while enhancing the localization accuracy.

We built and evaluated a proof-of-concept prototype of RL-SAR. Experimental evaluation demonstrates a mean 3D localization accuracy of 0.244m and the capability to locate 15 tags within an average scanning distance of 19.14 m. We compared our algorithm to naive baselines and show that the baselines require 86% longer trajectory than RL-SAR. Our results show the potential for achieving robust and efficient localization to enhance the current inventory processes across the manufacturing, retail, and logistics sectors.

Thesis Supervisor: Fadel Adib
Title: Associate Professor

Acknowledgments

First of all, I would like to thank Prof. Fadel Adib. I am grateful for your mentorship, your faith in me, and your consistent support over the past year. Thank you for taking the time to provide me with feedback and guidance on my thesis. I am also grateful for the friendship and help I received from everyone in Signal Kinetics. It was truly a pleasure to work with such a group of brilliant, fun individuals. I owe a special thanks to Isaac, Tara, Laura, and Purui for mentoring me and answering the many questions I had about the research. To Unsoo, thank you for your friendship and support. I have thoroughly enjoyed sharing an office with you, as well as all the interesting conversations we've had. I would also like to thank my family for raising me and providing continuous support over the past 24 years and throughout my education. Finally, I'm very grateful that the ceiling-mounted robot track has continued to function reliably over the past year. Regardless of how much I've abused it, it has never complained and has continued to perform consistently. The research is sponsored by the National Science Foundation (NSF), the Sloan Research Fellowship, Office of Naval Research (ONR), and the MIT Media Lab.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

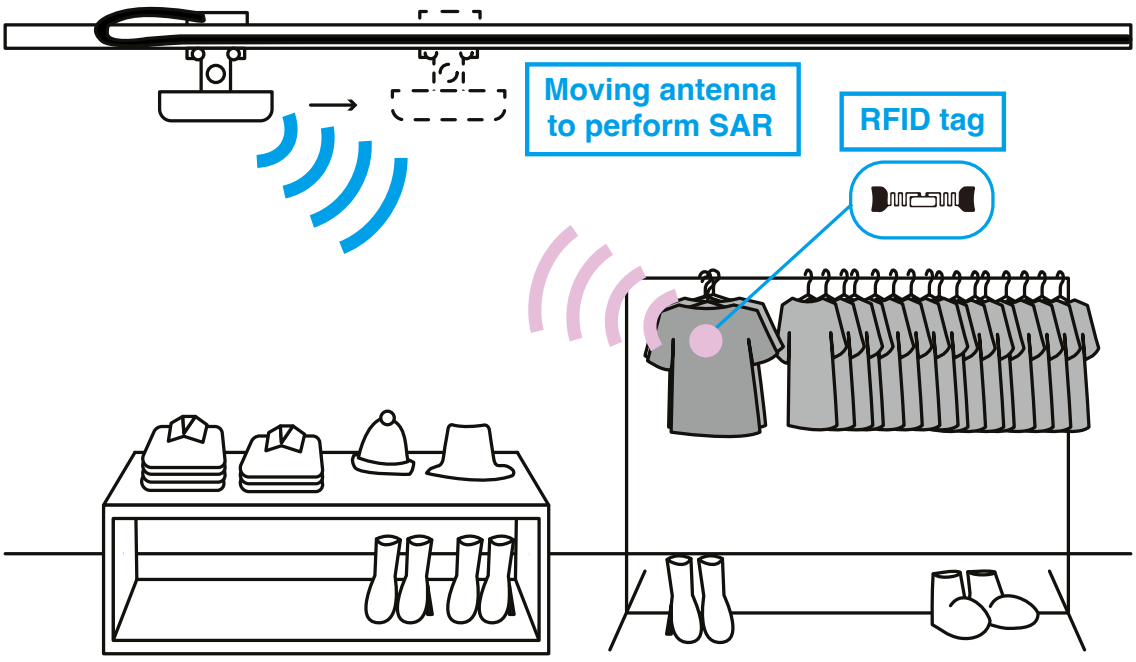

Figure 1-1: **RL-SAR in a store setting.** RL-SAR moves an ceiling-mounted antenna autonomously to perform Synthetic Aperture Radar (SAR) and localize RFID tags.

The adoption of Ultra High Frequency (UHF) RFID technology is expanding at an unprecedented rate in industries such as retail, manufacturing, and warehousing. In 2021, over 30 billion RFID tags were sold, with 70% used in apparel, footwear, and other retail products [1]. These tags are batteryless, small, and inexpensive (around 3

cents), which allows companies to apply the technology to all their products for more efficient inventory management. Many retailers use RFID systems to track inventory availability, perform automatic checkout, and detect theft activities. Apart from these applications, there is considerable interest in achieving fine-grained localization of these RFID-tagged items to enable precise inventory tracking. State-of-the-art RFID localization systems have demonstrated a centimeter level of accuracy using techniques such as Time-of-flight ranging [2] and Synthetic Aperture Radar (SAR) [3–8]. This level of accuracy paves the way for enabling us to differentiate between items on the same shelf and potentially create a digital map of items, which would be beneficial for locating specific items or managing inventory more efficiently.

Unfortunately, existing systems for fine-grained RFID localization in these environments are inefficient. One of the most popular approaches for performing such localization is to mount an RFID antenna on a mobile robot that moves on a predefined trajectory (e.g., traversing a dense spatial grid), collecting measurements from different locations in order to localize RFID tags in the environment [3, 4, 9–13]. However, such an approach suffers in scenarios where a subset of tags has not been localized accurately; this is because the robot either needs to repeat the entire trajectory for scanning again or it must settle for poor localization accuracy for these tags. To enable more accurate and efficient tag localization, some recent work [14] has considered using reinforcement learning to choose efficient locations for the robot to move to in order to collect RF measurements for localization; such an approach speeds up the task of localizing an individual RFID, but must be repeated for every tag in the environment, making it inefficient for localizing all tags in the environment.

In this thesis, we investigate whether we can achieve autonomous RFID localization more efficiently for multiple RFID tags using SAR with a trajectory optimization algorithm. SAR-based RFID localization leverages the relative movement between the antenna and the tag to create virtual antenna arrays and improves localization accuracy by increasing the aperture of the tag measurements. At a high level, our aim is to construct a robotic system that can optimally select a SAR trajectory to collect RF measurements and minimize the total scanning distance needed to localize all the

target tags in the environment. With this system, we can not only localize multiple tags more efficiently but also perform additional scans to adaptively relocalize tags with lower confidence.

Translating this concept into a practical system presents several challenges. First, we need to develop an autonomous robotic system that synchronizes antenna mobility with RF measurements to achieve real-time data collection. Also, this system employs Commercial Off-The-Shelf (COTS) hardware, which introduces the issue of phase ambiguity in RF measurements. Furthermore, the system needs to consider the localization confidence for all the tags and identify the locations to take new measurements. To address these challenges, we developed RL-SAR, an end-to-end autonomous SAR-based RFID localization system that operates based on an antenna moving on a ceiling-mounted 2D track. The core part of this system is a Reinforcement Learning (RL) framework that identifies the next best SAR trajectory for localizing multiple tags. We explore various confidence metrics and their relationship to the system's performance. The RL algorithm exploits the confidence metrics of the target tags and learns a policy to minimize the trajectory length while enhancing their localization confidence. Through this algorithm, RL-SAR can operate autonomously in the background of environments such as stores or warehouses, continuously tracking many RFID-tagged inventory items simultaneously, as shown in Figure 1-1.

I built and evaluated the system in a real-world scenario where RL-SAR conducted scans in an indoor environment mimicking a store setting and visualized the localization results on a 3D indoor map. This application also includes a depth camera to incorporate 3D point cloud data and achieve real-time scanning of the environment. This vision-based 3D data provides the system with the necessary information about the environment to enable the creation of a digital twin of the area. This represents the concept of an RFID localization system that can be used in the retail industry, creating a digital twin of the tagged items within the store.

Chapter 2

Background

In this section, we explore the associated studies on RFID localization and delve into state-of-the-art RFID localization systems with RF measurement path optimization. Moreover, we examine the necessary background knowledge and requirements for designing a SAR-based RFID localization system using Commercial Off-The-Shelf (COTS) hardware.

2.1 Related Work

We offer an overview of various RFID localization methods, drawing particular attention to the latest advancements in RFID localization systems with RF measurement path optimization.

2.1.1 RFID Localization

RFID localization can be achieved using active [15–17] and passive tags. In the realm of large-scale RFID-tagged item tracking, passive tags are frequently used due to their compact size and lower cost. Passive UHF RFID localization systems rely on measurements derived from the backscattered signal emitted by the tag. These systems can be broadly categorized into signal strength-based, time of flight-based, direction-of-arrival-based, and SAR-based methods.

Signal strength, also known as Received Signal Strength Indicator (RSSI), based approaches depend on the amplitude of the backscattered signal [18, 19]. However, the RSSI is highly sensitive to factors such as multipath propagation, attenuation due to obstructions, tag material, and orientation, which can result in localization errors of up to a meter.

Systems like RFind [2] and TurboTrack [20] leverage the emulated wide bandwidth and incorporate a super-resolution technique to achieve sub-centimeter level accuracy even in a multipath-rich environment. They sample the RF measurements over a wide bandwidth and apply the Inverse Fast Fourier Transform (IFFT) to recover the distance through time-of-flight. Unfortunately, this method requires a very wide range of emulated frequencies and needs specialized hardware for localization. Current COTS hardware is not suitable for this technique. Angle of arrival localization techniques capitalize on the phase difference between measurements from different antennas to estimate the direction of the backscattered signal [21]. However, scaling such a system poses challenges as it would necessitate many antennas to deliver adequate spatial resolution.

The Synthetic Aperture Radar (SAR) approach has been widely used in state-of-the-art RFID localization systems due to its ability to achieve good spatial resolution [3–5, 7, 11–13, 22–25]. These advanced systems often incorporate Visual Inertial Odometry (VIO) to collect a known trajectory of measurements and exploit wireless channel estimation to create a SAR hologram for localization. Existing SAR-based localizations can be primarily separated into two categories. The first category is human-motion-based SAR localization. Systems like XAR [26] combine an RF antenna with an augmented reality headset to enable SAR RFID localization based on human motion.

The second category revolves around robot-based SAR localization. Systems such as PinIt [3] use mobile robot-based SAR to extract the multipath profile of each tag and localize based on the nearest reference tag. Other researchers have utilized ground or flying robots to collect measurements at various locations, often relying on generating RF holograms to find the best match for the phase sequence calculated

based on prior knowledge of the robot's motion trajectory.

2.1.2 RFID Localization Systems with Path Optimization

Current methods for fine-grained RFID localization lack efficiency. A common strategy for achieving RFID localization involves attaching an RFID antenna to a mobile robot, which then follows a pre-defined trajectory [3, 4, 9–13]. This approach is not efficient when it comes to large amount of tags. If the system cannot find the tag, the robot has to go over the same path again. To improve the localization accuracy and efficient, some recent works employ path optimization to guide the user or robot to the optimal next locations for attempting to read RFID tags. RF-AR [27] explores the synergies between humans and the RFID localization system, introducing a path optimization algorithm based on Dilution of Precision (DoP) improvements. This optimal path is displayed as an arrow on the AR headset worn by the user. However, RF-AR only optimizes the trajectory for finding a single tag, and is inefficient when localizing multiple RFID tags simultaneously.

RFusion [14] uses a robotic arm to move the antenna and capture RF measurements from several locations. The system incorporates an RF-visual reinforcement learning algorithm to determine the optimal next vantage point to which the arm should move and take measurements. Guided by this RL-based algorithm, the robotic arm repeats the measurement process until it achieves sufficient localization confidence in its estimated location of the RFID tag. Unfortunately, much like RF-AR, RFusion is also solely designed to optimally search for a single RFID tag.

In this thesis, our objective is to construct a robotic system capable of moving the antenna to gather measurements and calculate location estimates using SAR. Furthermore, we aim to develop an RL-based path optimization algorithm that identifies the next best SAR trajectory for localizing multiple target tags.

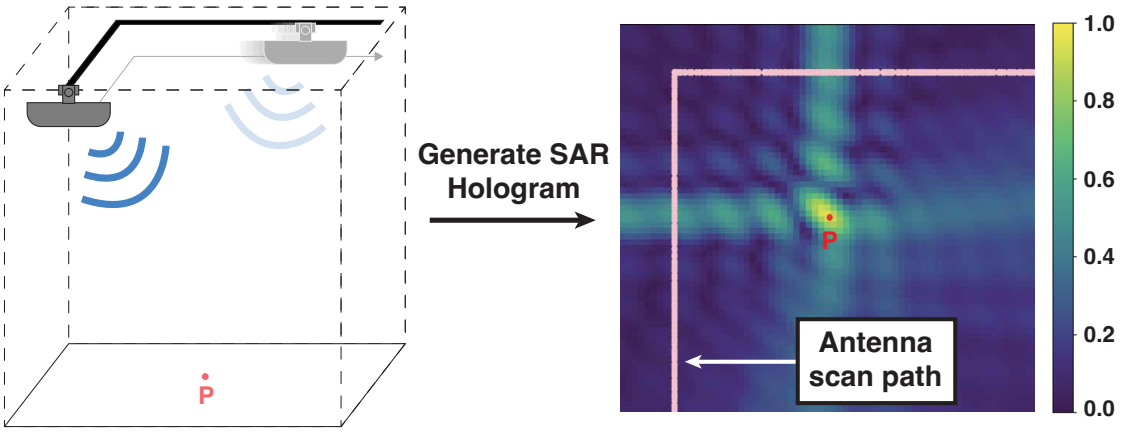

Figure 2-1: **Illustration of SAR.** The antenna moves along an L-shaped trajectory and collects wireless channel measurements at each point. These measurements are then used to create a SAR hologram, which is a likelihood map of the location of tag P. The hologram values are normalized between 0 and 1.

2.2 Primer on SAR-based RFID Localization

This section provides the fundamental background for SAR-based RFID localization system. First, we introduce the basic concept of SAR. Then, we explore the system requirements and constraints involved in designing a SAR-based RFID localization system using Commercial off-the-shelf (COTS) hardware.

2.2.1 SAR Localization

Synthetic Aperture Radar, or SAR, is a technique that enhances the spatial resolution of RF measurements by moving the antenna. In contrast to conventional antenna arrays, it offers the advantage of requiring fewer antennas to achieve the same spatial resolution. Broadly, SAR computes an RF hologram based on wireless channel estimations at each RFID measurement point. This RF hologram represents the likelihood of a tag being at each grid location in the space. The grid with the maximum value is estimated as the location most likely to contain that tag.

The SAR approach requires that the system estimates the current location of the antenna while it's moving. Existing systems often use Visual and/or Inertial Odometry (VIO) or stepper motor step counts to estimate the relative movement

between measurements [3, 27, 28]. In this thesis, RL-SAR leverages the stepper motor step counts provided by the ceiling-mounted 2D track to obtain the antenna position estimates required to perform SAR. Figure 2-1 illustrates the antenna moving in an L-shaped trajectory and creating a SAR hologram based on wireless channel estimation.

As the antenna moves on the ceiling-mounted track, it simultaneously transmits RF signals to communicate with the passive RFID tags in the environment. After the tag receives the signal from the antenna, it backscatters a signal with its identifier to the same antenna. RL-SAR exploits this received signal to estimate the wireless channel for measurement point i based on the following equation [29].

$$h_i = A_i e^{-j\phi_i} \tag{2.1}$$

where A_i is the measured amplitude and ϕ_i is the measured phase value at point i.

This wireless channel estimation is repeated at various measurement points, and the channel information can be combined to calculate the likelihood P at every point in the environment. Equation 2.2 provides a definition of the likelihood P at coordinate (x, y, z).

$$P(x, y, z) = \left| \frac{1}{N} \sum_{i=1}^{N} h_i e^{j\frac{2\pi d_i(x,y,z)}{\lambda}} \right| \tag{2.2}$$

where N is the total number of measurements, h_i is the channel estimation of the i^{th} location, d_i is the round trip distance from (x, y, z) to the i^{th} location, and λ is the wavelength of the received signal. Finally, the tag location p_{est} can be estimated by the following equation.

$$p_{est} = \texttt{argmax}_{(x,y,z)}(P(x, y, z)) \tag{2.3}$$

2.2.2 Measurement Interval Requirements

When performing SAR localization, the intervals between measurements are constrained by a value to avoid multiple candidates on the SAR hologram. This inter-measurement constraint is defined to eliminate the spatial aliasing effect caused by grating lobes [30]. If the measurement locations are separated by more than $\lambda/4$, there will be multiple candidates for each possible on the SAR hologram that result in the same likelihood value. This is non-ideal for locating a large number of tags and compromises the localization accuracy of our system. Therefore, RL-SAR incorporates a filter to ensure that subsequent measurements in the same trajectory are within this maximum interval limit.

2.2.3 Constraints on COTS Hardware

According to the COTS hardware - ThingMagic M6e RFID reader [31] that we used and the report from this research [32], these COTS RFID readers only return $0 - \pi$ phase values instead of $0 - 2\pi$ phase values. This creates a π ambiguity for the system and causes problems during SAR computation. Firstly, the maximum gap derived from §2.2.2 becomes one-half of the original value due to the fact that the effective wavelength is half of the original wavelength.

The maximum gap constraints now become $\lambda/8$ when using COTS hardware. Furthermore, to aggregate the wireless channel estimation when calculating the RF hologram, we need to perform signal phase unwrapping as described in §3.2.2.

Chapter 3

RL-based SAR RFID Localization

This chapter first provides an overview of the RL-SAR system, followed by a detailed description of the autonomous SAR-based localization pipeline. Lastly, we introduce the RL-based trajectory optimization algorithm designed to identify tags with lower confidence and suggest efficient scanning trajectories for their localization.

3.1 System Overview

RL-SAR is an autonomous platform for SAR-based RFID localization using a 2D robotic track system. It is installed on the ceiling with an RFID antenna installed on the track. The antenna can move in both X and Y direction to collect measurements from various locations. These measurements are utilized to perform SAR calculations and localize the tag. Figure 3-1 provides an overview of the RL-SAR system.

The localization procedure starts with the RL-SAR controller sending preset trajectory movement commands to the 2D ceiling-mounted track system. As it starts moving, it also simultaneously enables the RFID reader to collect measurements along the way. After RL-SAR finishes one scan, the data is sent to the local server for preprocessing such as timestamp synchronization, phase unwrapping, and maximum gap filtering. The processed data is then used to compute the RF SAR hologram, which is aggregated with previous measurements to form a cumulative RF hologram. From this, we are able to derive the estimated location of the tag and calculate its con-

fidence intervals. If the confidence intervals satisfy the predetermined convergence criteria, the system returns a localization result. If they do not, it activates the adaptive relocalization procedure, and RL-SAR recommends an optimal path for the next scan.

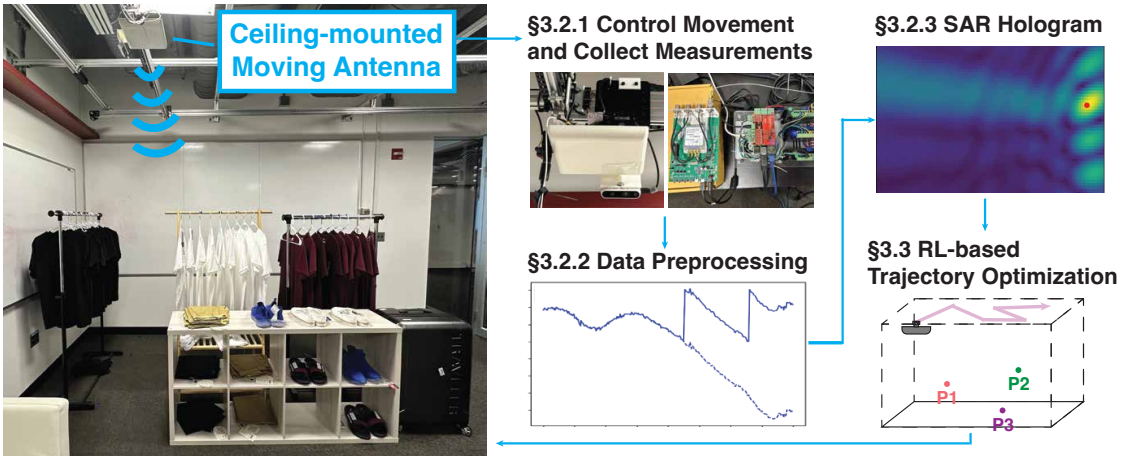

Figure 3-1: **RL-SAR Overview.** RL-SAR utilizes a ceiling-mounted moving antenna to perform SAR and localize RFID tags. §3.2.1 describes the process of controlling the movement and collect RF measurements. §3.2.2 describes the data preprocessing steps. §3.2.3 describes the SAR RF hologram computation. §3.3 shows how to use RL-based trajectory optimization algorithm to find the optimal trajectory for doing SAR on finding multiple tags with low confidence.

3.2 End to End RFID Localization using SAR

In the section, we describe how RL-SAR achieves autonomous SAR based 3D localization in detail, including how it controls movement, collects measurements, performs preprocessing, computes SAR hologram, and calculates confidence metrics.

3.2.1 Control Movement and Collect Measurements

To achieve SAR, it is necessary to control the movement of the antenna and know its position at each measurement point. We define the basic movement at index i as (x_i, y_i, s_i, r_i), where x_i and y_i are the next position where we want the antenna to move to, s_i defines the speed that the antenna is moving at, and r_i defines whether we want

the RFID reader to read during this movement. When we controls the movement of RL-SAR's antenna, we provide a sequence of basic movements, or a trajectory:

$$\tau = [(x_0, y_0, s_0, r_0), (x_1, y_1, s_1, r_1), \ldots, (x_i, y_i, s_i, r_i)] \tag{3.1}$$

where each of the x_i, y_i coordinates in τ represents a setpoint in a trajectory, and RL-SAR continues to read the current location of the antenna to see if it arrives at the setpoint before executing another movement toward the next setpoint. Note that x_i and y_i are in meters, and s_i has the unit of mm/min. Also, the coordinate (x_i, y_i) needs to be constrained in the (x_{min}, x_{max}), (y_{min}, y_{max}) bound. [1]

During the movement, RL-SAR continues to request current coordinate feedback from the linear track system and save the feedback information as (x_n, y_n, t_n). (x_n, y_n) is the coordinate at which the feedback is collected from, and t_n is the timestamp information in epoch time. After scanning through a trajectory, we get a list of feedback positions Ω at different timestamps:

$$\Omega = [(x_0, y_0, t_0), (x_1, y_1, t_1), ..., (x_n, y_n, t_n)] \tag{3.2}$$

If the read flag, r_i, is enabled in the movement, RL-SAR controls the RFID reader to simultaneously read RFID tags in the environment. RL-SAR communicates with the RFID reader through the Mercury API [33] to begin reading the tags, and the reader responds when a successful measurement has been collected. Each measurement at time t_k contains $(epc_k, rssi_k, \phi_k, f_k, t_k)$ where epc_k is the RFID tag electronic product code (EPC) identifier, $rssi_k$ is the RSSI of the received signal, ϕ_k is the measured phase value, and f_k is the frequency of the received signal. After RL-SAR completes a scan, we can get a list of measurements collected by the RFID reader:

$$\Lambda = [(epc_0, rssi_0, \phi_0, f_0, t_0), (epc_1, rssi_1, \phi_1, f_1, t_1), ..., (epc_k, rssi_k, \phi_k, f_k, t_k)] \tag{3.3}$$

[1] RL-SAR uses $(0, 3)$ for (x_{min}, x_{max}), and (0, 2) for (y_{min}, y_{max}). It defines an area of 3m x 2m that the antenna can move in.

At the beginning of the autonomous SAR localization task, RL-SAR moves along a preset trajectory τ_{init} to maximize the possibility of reading all the tags in the environment. This preset trajectory is a rectangular trajectory that has a 0.5m offset from all four bounds of the ceiling-mounted track system. [2]

3.2.2 Data Preprocessing

To prepare data for computing the SAR RF hologram, RL-SAR performs data preprocessing tasks which include target tags filtering, timestamps synchronization, maximum gap filtering, and phase unwrapping. In this section, we discuss these processes in detail in the following subsections.

Target Tags Filtering

Before doing the following preprocessing tasks, we define the target tags that we aim to localize. In this thesis, we define a set of target tags S_t before the localization process starts. In actual application, this set may be defined during the previous inventory round, with known tags' EPC corresponding to each of the items in the database. Nonetheless, if we want to search and localize all tags in the target area, RL-SAR also works with S_t containing all the unique RFID tag EPCs that are in the measurement list Λ.

For all of the measurements in Λ, we categorize the measurements by their EPC identifiers. Therefore, we formally define the measurement list for tag m as:

$$\Lambda_m = [..., (epc_k, rssi_k, \phi_k, f_k, t_k), ...] \qquad \forall\ epc_k = \text{EPC of tag } m \tag{3.4}$$

where Λ_m filters out all the measurements of tag m from Λ using the RFID tag's EPC. Note that tag m should be a target tag that is included in S_t.

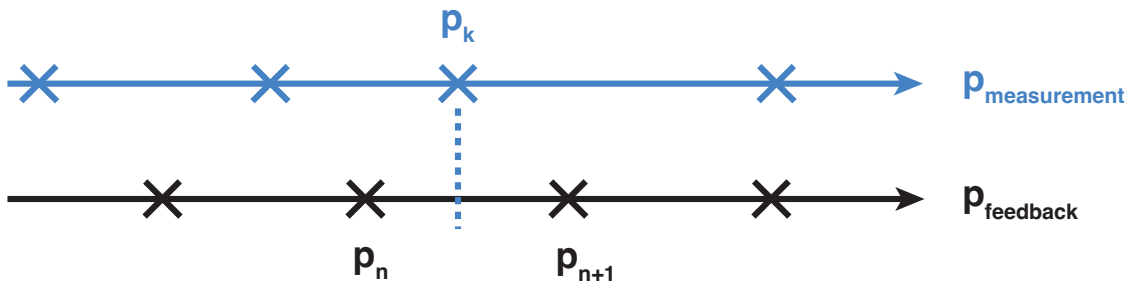

Figure 3-2: **Synchronizing RF Measurements with Position Feedback.** The timestamp of the RF measurement and the timestamp of the position feedback are not synchronized. $p_{measurement}$ represents the positions on the track that the measurement was collected, and $p_{feedback}$ is the periodic position feedback from the track. We use linear interpolation on the position feedback data (p_n and p_{n+1}) to find the exact position p_k, when the measurement was collected.

Timestamps Synchronization

Recall from equation 2.2 that we need the exact position of each measurement to compute SAR RF hologram. However, there might not be an exact same timestamp in Ω for the measurements in Λ_m to find their corresponding locations. As we can see from Figure 3-2, both the location feedback and the measurements are taken at different timestamps, and we aim to find the location p_k where RL-SAR collects the k^{th} measurement.

As the antenna moves at a constant speed on the 2D ceiling-mounted track, we can apply linear interpolation to find precisely where the measurements in Λ_m were collected from. For each measurement timestamp t_k, we can find the two closest position feedback points in Ω as (x_n, y_n, t_n) and $(x_{n+1}, y_{n+1}, t_{n+1})$. The location p_k of the measurement k can be calculated by the following equation:

$$p_k = \left(1 - \frac{t_k - t_n}{t_{n+1} - t_n}\right) \cdot p_n + \frac{t_k - t_n}{t_{n+1} - t_n} \cdot p_{n+1} \tag{3.5}$$

And the synchronized RF measurement data can now include the position data

[2]Formally, τ_{init} is defined as $[(0.5, 0.5, 1000, 1), (0.5, 1.5, 1000, 1), (2.5, 1.5, 1000, 1), (2.5, 0.5, 1000, 1), (0.5, 0.5, 1000, 1)]$

and transform the measurement list to become:

$$\Lambda_m{}' = [..., (epc_k, rssi_k, \phi_k, f_k, x_k, y_k, t_k), ...] \tag{3.6}$$

Maximum Gap Filtering

Recall from §2.2.2 that there is a requirement in antenna array measurement spacing to avoid multiple candidates in SAR localization. For neighboring measurements, we ensure their locations are separated by a gap smaller than $\lambda/8$. Therefore, we filter $\Lambda_m{}'$ based on the interval requirement and group measurements into different clusters s. All the measurements within cluster s are separated no more than $\lambda/8$ from their neighbors. We define the measurement list of cluster s as $\Lambda_{ms}{}'$. For each $\Lambda_{ms}{}'$, we also ensure that there are at least 5 measurements in the list to provide a meaningful update to SAR computation. We discard all $\Lambda_{ms}{}'$ with a measurement count of less than 5.

Phase Unwrapping

Recall from §2.2.3 that because RL-SAR uses a COTS RFID reader for collecting measurements, there exists a π ambiguity in the collected phase values. Therefore, we apply a phase unwrapping algorithm to unwrap the phase to prevent phase ambiguity from leading to faulty results. As we can see from Figure 3-3, the standard blue line shows the raw phase reading collected by the RFID reader, and it has several discontinuities that jump from 0 to π. The phase unwrapping algorithm looks for the sudden jump and adds an offset to the raw phase reading in order to eliminate the discontinuities. The unwrapped phase ϕ'_k can be formally defined as follows:

$$\phi'_k = \begin{cases} \phi_k + \pi & \text{if } \phi_n - \phi_{n-1} \leq -\frac{\pi}{2} \\ \phi_k - \pi & \text{if } \phi_n - \phi_{n-1} \geq \frac{\pi}{2} \\ \phi_k & \text{otherwise} \end{cases} \tag{3.7}$$

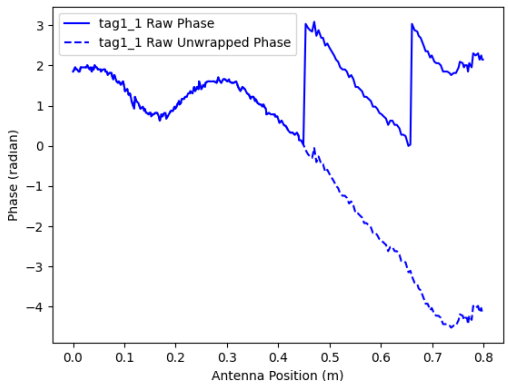

Figure 3-3: **Raw and Unwrapped Phase vs Antenna Position.** The antenna is controlled to move only in the y axis direction. Therefore, the antenna position represents the y coordinate. The standard blue line shows the raw phase readings from the RFID reader, while the dotted line shows the unwrapped phase. We can see that the raw phase readings wrap around at $[0, \pi]$ and RL-SAR unwraps the phase whenever there's a sudden jump from 0 to π.

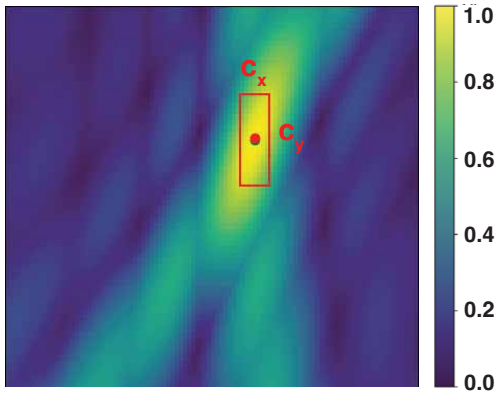

Figure 3-4: **SAR Hologram with Confidence Interval.** The SAR RF hologram shows the likelihood value of having the target tag at each grid in the 3D space. The confidence interval is defined by c_x and c_y, the length and width of the bounding box that encompass locations that are likely to have the tag. Smaller confidence intervals indicate lower variance in the localization result, which implies higher confidence in the estimated location.

It means that if the algorithm senses a discontinuity of greater than $\pi/2$, it will add or subtract an offset π to unwrap the raw phase values. Using this algorithm, we find phase discontinuities in all ${\Lambda_{ms}}'$ and update the phase value for every measurement in the list accordingly.

3.2.3 SAR RF Hologram

Based on the preprocessed measurement data ${\Lambda_{ms}}'$, we compute the SAR RF Hologram for each tag m. RL-SAR defines the localization area with bounds: (x_{min}, x_{max}), (y_{min}, y_{max}), and (z_{min}, z_{max}) [3]. To perform SAR RF hologram, we discretize the 3D space and calculate the likelihood of every grid. The grid size ($x_{grid} = 0.1m, y_{grid} = 0.1m, z_{grid} = 0.2m$) is chosen so that the SAR computation can be completed in a reasonable timeframe with an acceptable accuracy. The likelihood value for each grid in the hologram is calculated based on the equation 2.2. However, when we perform

[3]RL-SAR uses $(0, 3)$ for (x_{min}, x_{max}), $(0, 2)$ for (y_{min}, y_{max}), and $(0, 1.6)$ for (z_{min}, z_{max})

wireless channel estimation based on the measured data, the measured RSSI values can be easily affected by multipath or obstructions in the environment, leading to inaccurate channel estimation. Therefore, we set all the amplitude of all channel estimation to 1 and only use the phase information to calculate the likelihood value. Formally, the likelihood RF hologram for the measurement list ${\Lambda_{ms}}'$ is defined as:

$$P_{ms}(x, y, z) = \left| \frac{1}{N} \sum_{i=1}^{N} e^{-j\phi_i} e^{j\frac{2\pi d_i(x,y,z)}{\lambda}} \right| \tag{3.8}$$

For each tag m, we have multiple measurements ${\Lambda_{ms}}'$ as the process continues and perform more scans. In order to combine these results, we aggregate the hologram values calculated with a single measurement list and update the global SAR RF hologram for each tag. In order to reinforce the RF hologram with quality measurements that result in a higher likelihood value, we apply a weighted number for each hologram before aggregating the hologram values to the global hologram. The weighted number is chosen to be the maximum likelihood of that particular hologram, and the global SAR RF hologram for tag m can be formally defined as:

$$P_m(x, y, z) = \sum_{s} P_{ms}(x, y, z) \cdot \max_{(x,y,z)} P_{ms}(x, y, z) \tag{3.9}$$

Based on the global SAR RF hologram, we find the coordinate of the grid that contains the maximum likelihood value and output the estimated location $\hat{p_m}$ for the specific tag m.

$$\hat{p_m} = \texttt{argmax}_{(x,y,z)}(P_m(x, y, z)) \tag{3.10}$$

3.2.4 Confidence Metrics

In this section, we discuss different metrics that represent the confidence level for the location estimate $\hat{p_m}$. Specifically, we compute the confidence interval and the Dilution of Precision (DoP) of the tag's location estimate based on the current measurements.

Confidence Interval

To understand the confidence level across three different axes, we define the confidence interval as a region that encompasses all (x, y, z) grid locations where the likelihood value is greater than -0.75 dB of the maximum likelihood value. The confidence interval includes three numbers: (c_x, c_y, c_z), and it describes the maximum span in all three directions that include all grid locations meeting the likelihood requirement of -0.75 dB. The larger the confidence interval, the greater the variance of the localization results, indicating a lower localization confidence. Conversely, smaller confidence interval values indicate higher localization confidence. Figure 3-4 displays the SAR RF Hologram and the confidence interval of this localization result.

Dilution of Precision

Dilution of Precision (DoP) is a concept that describes how errors in the measurements will affect the final state estimation [34]. It is widely used in satellite navigation like GPS and can be used to express the level of uncertainty based on the geometry of the measurement points and the estimated target location. When all measurements are from approximately the same regions, the DoP value will be large, and the error in the measurement creates a larger uncertainty in the estimated location. Conversely, if the measurements are taken from a larger variety of locations, DoP will be smaller. We can formally calculate the DoP value based on measurement positions and the estimated tag position as follows [4]:

$$A = \begin{bmatrix} \frac{x_1 - \hat{x_m}}{R_1} & \frac{y_1 - \hat{y_m}}{R_1} \\ \frac{x_2 - \hat{x_m}}{R_2} & \frac{y_2 - \hat{y_m}}{R_2} \\ \vdots & \vdots \\ \frac{x_k - \hat{x_m}}{R_k} & \frac{y_k - \hat{y_m}}{R_k} \end{bmatrix} \tag{3.11}$$

[4] In RL-SAR, we only focus on DoP values in the x and y directions, as our robotic system has no freedom in the z direction to enhance DoP.

$$Q = \left(A^T A\right)^{-1}, \quad \text{DOP} = \sqrt{tr(Q)} \tag{3.12}$$

where (x_k, y_k) correspond to the location of the antenna for measurement k, $(\hat{x_m}, \hat{y_m})$ is the estimated target RFID tag location calculated in equation 3.10, R_k is the distance from the measurement location to the estimated tag location, and $tr(.)$ is the trace of matrix.

RL-SAR uses confidence interval and DoP as indicators to suggest whether the tag location estimate has converged and has high possibility of meeting the localization accuracy requirements. The convergence criteria is defined as follow:

$$\begin{cases} c_x < 0.3, \\ c_y < 0.3, \\ c_z < 0.6, \\ \text{DOP} < 0.008 \end{cases} \tag{3.13}$$

When the above criteria holds, the system determines that a specific tag m is localized successfully. Details on the microbenchmarks of these values can be found in §5.1, where we compare different confidence intervals and DoP to the localization accuracy.

3.3 RL-based Trajectory Optimization

In this section, we detail RL-SAR's approach to determine next best scanning trajectory and relocalize tags with low localization confidence. Recal from §3.2.1 that RL-SAR initiates the localization process by sending τ_{init} to the system. After this initial scan, if the confidence interval and DoP do not meet the convergence criteria 3.13, RL-SAR begins to identify the next location for a scan. The goal of this algorithm is to achieve convergence criteria for all the tags while minimizing the antenna scanning distance.

Given the estimated locations of the tags, one might assume that the optimal path would be to set all these candidate locations as the setpoints of the trajectories and

move toward all of them. However, this strategy does not account for the confidence interval, DoP, or whether it can optimize multiple tags simultaneously. To address these confidence metrics and the ability to optimize localization for multiple tags concurrently, we formulate this as an optimization problem aimed at minimizing the scanning distance for RL-SAR's antenna. We then employ a reinforcement learning network to solve this problem, enabling the reinforcement learning agent to learn an optimal policy in the simulator. The simulator setup is detailed in §4.2.

At a high level, the reinforcement learning network attempts to recommend the next best location for the antenna to move to and collect RFID tag readings along the way. It strives to minimize the total scanning distance while successfully localizing all the tags. Let N denote the number of tags we aim to optimize using the RL algorithm [5]. We assign 5 state parameters for each of the tags as: $(x_{est}, y_{est}, c_x, c_y, L_f)$. Here, (x_{est}, y_{est}) represents the current estimated location of a specific tag, (c_x, c_y) indicates the confidence interval of the localization results, and L_f is a flag that records whether the tag has been localized. To simplify the complexity of the reinforcement learning problem, we only chose the confidence interval as the metric to include in the state representation. It's important to note that although we fix the number of tags to N in our RL algorithm, the method can locate an arbitrary number of tags by feeding a subset of N tags into the RL network. In RL-SAR, we select the closest N tags to the current antenna location to iteratively optimize the localization of the tags at a closer distance. In addition, we have 3 fixed parameters: (s_{total}, a_x, a_y), which represent the total scanning distance, current antenna position x, and current antenna position y. Consequently, our state space has a size of $5N + 3$, and the representation S can be defined as follows:

$$S = \{(x^i_{est},\, y^i_{est},\, c^i_x,\, c^i_y,\, L^i_f)_{i=1}^{N},\ s_{total},\ a_x,\ a_y\} \tag{3.14}$$

The RL network takes in the state S and outputs an action that represents the next best coordinate that it suggests RL-SAR to move to. The action space has two

[5] RL-SAR sets $N = 5$

values: $AC = (n_x, n_y)$, the next x and y coordinate.

As the RL agent performs a scan based on the output action in the simulator, the environment calculates a reward. The reward is primarily dependent on the total traveled distance of the scanning antenna. The agent will also be rewarded if it localizes all the tags. However, it includes a penalty term when the movement is too small or doesn't localize all the tag as it reaches maximum allowed steps. The reward function is defined as follow[6]:

$$R = -\alpha \cdot s_{total} + \beta \cdot ac_{penalty} \tag{3.15}$$

where $ac_{penalty}$ penalizes antenna scan length shorter than 0.5m. We define the antenna scan length as $\Delta AC = ||AC_i - AC_{i-1}||$, the length between coordinate at time i and time $i-1$. Formally, this term is defined as:

$$ac_{penalty} = \min\left(-\left(1 - \frac{\Delta AC}{0.5}\right), 0\right) \tag{3.16}$$

Based on the state, action, and reward definitions, we apply the state-of-the-art RL policy optimization algorithm - Proximal Policy Optimization (PPO) to learn the optimal policy and minimize scanning trajectory length in order to localize all the tags [35]. PPO is a policy gradient method that incorporates a clipping function to prevent the improvement step from deviating too far, which could lead to performance collapse. It calculates the reward value based on interaction with the environment and computes an advantage estimate using the current value function. With these values, the PPO algorithm updates the network parameters by maximizing the PPO objective. Formally, the objective and update step can be defined as follows:

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right] \tag{3.17}$$

$$\theta_{t+1} = \texttt{argmax}_\theta \; L^{CLIP}(\theta) \tag{3.18}$$

where θ is the policy network parameters, $\hat{A}_t$ is the estimated advantage at time t,

[6]We set $\alpha = \frac{1}{180}$ and $\beta = 1$ to normalize both terms to $0-1$

$r_t(\theta)$ is the ratio of the probability of the action under the new policy to the probability under the old policy: $r_t(\theta) = \frac{\pi\theta(a_t|s_t)}{\pi\theta_{old}(a_t|s_t)}$, and ϵ is a hyperparameter that defines the maximum allowable change in policy, usually set to 0.2.

As the RL agent learns the policy in a simulation environment, the learned policy might not generalize well when implemented on the real hardware. In this thesis, we apply the system identification technique to attempt to close the gap between the simulation and the real-world environment. We experimented with various sets of values for system parameters such as the noise range of the measured phase value and the tag readability parameters, and evaluated the system performance on localization accuracy and convergence time based on the same scenario. Then, we selected a set of parameters that yields performance most similar to the real-world RL-SAR system. Detailed descriptions of the parameters and the values we used in this thesis are specified in §4.2. This approach narrows the gap between the simulation environment and the real-world system, enabling the RL agent to learn in the simulated world and apply the same policy to the real system.

Chapter 4

Implementation and Evaluation

In this section, we discuss the hardware and experimental setup for RL-SAR, the simulator setup for the RL agent to learn the optimal trajectory policy, and the evaluation metrics for the system.

4.1 Hardware and Experimental Setup

The RL-SAR setup includes a 2D ceiling-mounted track system, a circularly polarized antenna attached to the ceiling-mounted movement platform, motor controller, a Thingmagic RFID reader, and a local server for computation. These components can be seen in Figure 4-1.

The 2D motion track system is constructed with two V-slot linear rails allocated to the X and Y axes respectively. These linear rails are attached to several T-slots, which are fastened to the wall near the room's ceiling. The system incorporates a NEMA 23 stepper motor [36] for each axis to pull the belt attached to a pulley, thus creating X and Y linear movement. The stepper motors are connected to the Protoneer Raspberry PI CNC board. This CNC board, equipped with a microcontroller running grbl [37], a motion control software for CNC machines, is mounted on top of a Raspberry PI 4 Model B and communicates through a serial port. We control the movement of the 2D motion track system by sending G code commands to the CNC board. The code that dispatches the G code via serial is implemented in Python. We

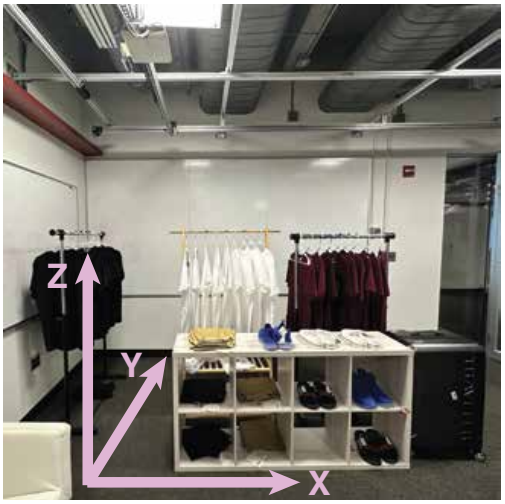

(a) Experimental Setup

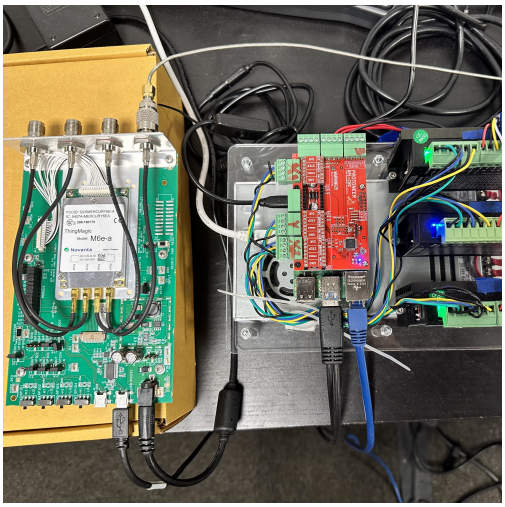

(b) Motor controller and RFID reader

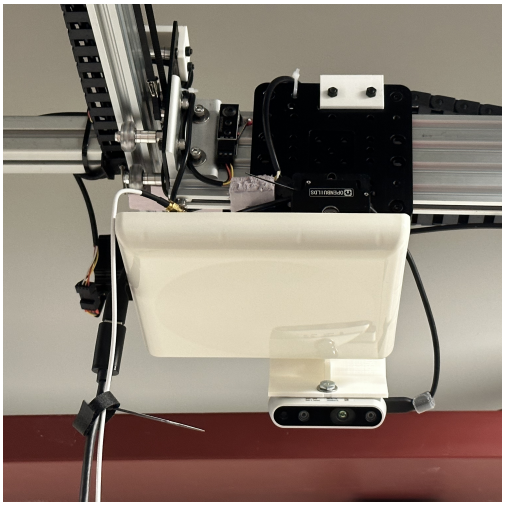

(c) Antenna for RF measurements

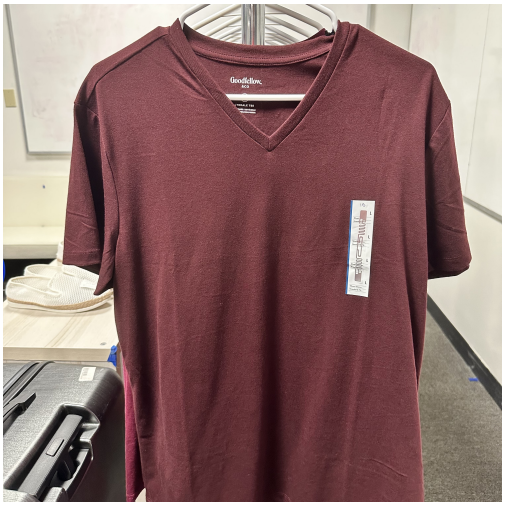

(d) Example RFID tagged item

Figure 4-1: **Experimental Setup for RL-SAR**

attached the MTI MT-242025/TRH/A (RHCP) Outdoor RFID Antenna [38] to the movement platform. In addition, we have installed the Intel Realsense D435i depth camera on the movement platform for constructing end-to-end digital twin applications and combining both sensing modalities to create a 3D map of the environment. Further details on this digital twin application are described in §7.

The mobile antenna is connected to the Thingmagic M6e UHF RFID reader [31], which links to the Raspberry PI via a USB cable. We control the Thingmagic RFID reader through the Mercury API [33]. It should be noted that we control the reader in such a way that it only transmits signals at the frequency of 915MHz, which falls

within the ISM band for UHF RFID applications. The implementation of the Mercury API and the program to collect RF measurements are written in Python code. The collected RF measurements and 2D motion track location feedback are transmitted to the local server via a ZMQ socket [39]. This initiates the data preprocessing and SAR RF hologram computation required to calculate the localization results. All of this code is written in Python and runs on a local server—an Ubuntu 22.04 machine equipped with a 12th Gen Intel(R) Core(TM) i9-12900K CPU @ 5.30GHz.

The experimental environment for RL-SAR is established in a multipath-rich environment featuring shelves, sofas, and clothing racks, simulating a standard store environment. The target items are labeled using standard off-the-shelf UHF RFID tags on their surfaces. The global origin of the environment is set at the ground level beneath the origin of the 2D motion track system. The coordinate system is established as a right-hand coordinate system, with z+ pointing upward, x+ to the right, and y+ to the front, as illustrated in Figure 4-1a.

4.2 Simulator Setup for RL training

We have developed a simulator for RL-SAR with the same dimensions of (x_{min}, x_{max}) and (y_{min}, y_{max}). To make the system interchangeable with the actual hardware setup, the simulator accepts the same trajectory input and outputs data in the same measurement format as the RL-SAR hardware. When the simulator receives the trajectory input, it simulates the stochastic nature of the measurement points and generates random read points along the trajectory. The simulator also integrates a tag readability parameter - t_r. This tag readability parameter is employed to determine the probability of the reader being able to read the tag at a particular point. It is set to reflect the real-world situation where the further the reader is from the tag, the lower the likelihood that the reader can read the tag's response. Consequently, the simulated RFID reader has a probability of t_r to read a tag, and this value is dependent on the simulated RSSI of the signal received by the reader. It can be

formally defined as follows:

$$t_r = \begin{cases} 1 - \frac{-48-RSSI}{-48-(-80)} & \text{if } -80 \leq RSSI \leq -48 \\ 1 & \text{if } RSSI > -48 \\ 0 & \text{if } RSSI < -80 \end{cases} \tag{4.1}$$

The simulated RL-SAR can read all the tags if the simulated RSSI value is greater than -48, and it cannot collect any measurement when the RSSI value is lower than -80. The simulated RSSI value is calculated based on a standard backscatter received signal link budget [40], and we add noise to the ideal distance, which is discussed next.

Regarding the simulated RF measurements for the output, the simulator computes the ideal RSSI and phase values based on the distance from the measurement point to the tag location. However, in the real-world RFID reader, the measurements can be distorted due to noise or obstructions in the environment. Therefore, we introduce a uniform random noise with the bound of $(-n_r, n_r)$ to the distance estimation, which affects both the RSSI and phase measurements. In this thesis, we choose $n_r = 0.015$. We have confirmed that our selection of n_r and definition of t_r yield localization accuracy and total scanning distance to convergence similar to those of the actual system when using the same scanning trajectories and the same localization method.

Our RL agent was trained in the simulator described above. The PPO algorithm was implemented using the stable baselines3 API [41]. The Adam optimizer was employed for training the policy. We used actor-critic as the main policy framework for training the RL agent. The minibatch size was set to 64. The network was trained over 150,000 interactions. The training process was performed on a machine running Ubuntu 22.04, equipped with an RTX 3090 graphics card, an Intel Core i9 CPU, and 32 GB DDR5 RAM.

	LoS	NLoS
Counts	11	4

	Clothes	Non-Clothes
Counts	5	10

Table 4.1: **Number of RFID Tagged Items by Category.** The experimental setup for RL-SAR includes 15 RFID-tagged items, randomly placed and categorized in the table. Non-Line-of-Sight (NLoS) tags represent RFID-tagged items that are hidden in shelf compartments, and cannot be directly seen by the antenna. Examples of non-clothes tagged items include shoes and small plastic boxes.

4.3 Evaluation

We evaluated RL-SAR in a multipath rich indoor environment designed to mimic a store setting. Figure 4-1a illustrates the evaluation environment for RL-SAR. This indoor area features common store furniture such as shelves, a sofa, and clothing racks. We selected clothes and shoes as the items we wanted to localize and attached standard UHF RFID tags to them. The shelves, racks, and sofa were randomly placed within the target localization area of RL-SAR, and the clothes and shoes were likewise randomly positioned on these furnitures in the area. Note that some of the tags were in non-line-of-sight situations due to obstructions like shelves or other items. Table 4.1 shows the profile for these randomly tagged items.

Before evaluating the system, we collected the target tags' EPC ids and input this information as the target tag set S_t, as described in §3.2.2. We also recorded the ground truth locations for these target tags relative to the origin of the environment. RL-SAR was evaluated based on two primary metrics: localization error and total scanning distance. The localization error is defined as the difference between the estimated tag location output by RL-SAR and the actual RFID tag location. The total scanning distance refers to the cumulative distance that the antenna traversed on the 2D ceiling mounted track to successfully locate all of the tags assigned in the target tag set S_t. Our goal is to find an optimal strategy that minimizes both the localization error and the total scanning distance.

Regarding RL-SAR's RL-based trajectory optimization strategy, we evaluated our policy alongside two different naive strategies to assess whether it demonstrated an improvement in minimizing the total scanning distance. The two naive strategies are the random trajectory and the exhaustive search strategy. The random trajectory

strategy generates random (x,y) coordinates as the next movement location for the antenna. The exhaustive search strategy, meanwhile, controls the robot to move in a rectangular spiral pattern and aims to cover all areas through a fixed pattern to exhaustively search the entire environment. This strategy recommends a path that starts from the border of the RL-SAR movement space and follows a rectangular spiral path inward towards the center of the region. The antenna begins by following the outermost four edges of the search space, offsetting the rectangular path inward by 0.1m at each iteration. If RL-SAR has not localized all the tags after reaching the center of the space, it continues its movement by moving outward again, hoping to gather more measurements along the way. While this naive strategy guarantees that the antenna reaches all locations in the search region, it does not optimize the total antenna scanning distance by effectively leveraging existing information.

Chapter 5

Microbenchmarks

5.1 Confidence Metrics

In our first microbenchmark, we evaluated various confidence metrics to demonstrate their relationship with the localization results. These metrics include confidence intervals, DoP, and aperture. While the first two metrics are explained in §3.2.4, the aperture refers to the span of the measurement points. For example, we define X aperture as the difference between the maximum x location and the minimum x location of all measurements.

In this experiment, we placed 5 Line-of-Sight tags in the environment, and implemented a random trajectory strategy with RL-SAR to create a variety of localization confidence profiles. Throughout the scan, we recorded each confidence metric at each step as well as the current location estimates. When the localization converged, we initiated a new round with new random scans. We repeated the process until we had collected data from 100 experiments. Figure 5-1 plots the localization error across three different confidence metrics. From the results, we make the following remarks:

- The localization error increases as DoP and confidence intervals increase, whereas the error decreases when the aperture of the measurement points increases.
- RL-SAR can consistently locate the tags when the DoP and confidence intervals are sufficiently low. In Section 3.2.4, we define the convergence criteria to be

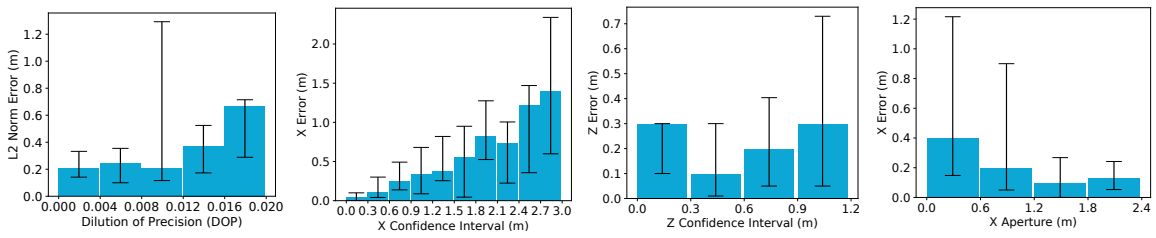

(a) DoP (b) X Confidence (c) Z Confidence (d) X Aperture

Figure 5-1: **Localization Error vs Various Confidence Metrics.** Each bar represents the median localization error of that bin. The error bar of each bar marks the 10^{th} percentile and 90^{th} percentile of the data. We only evaluated the aperture and confidence interval metrics for the X axis and not the Y axis because RL-SAR moves in XY direction, and these two axes are interchangeable.

DoP < 0.008 and confidence intervals $<$ (0.3, 0.3, 0.6). Figure 5-1a shows that the 3D median localization is less than 0.241m when DoP < 0.008. Additionally, when the confidence intervals are less than (0.3, 0.3, 0.6), Figure 5-1b shows an X error of 0.05m, while Figure 5-1c displays a median Z error of less than 0.3m. As the X and Y-axis performances are interchangeable, we can combine the error in three directions and estimate a 3D localization error of 0.308m. When all convergence criteria are met, we can expect that RL-SAR has a median 3D localization error of less than 0.308m.

This microbenchmark demonstrates that the confidence metrics we selected for RL-SAR can be used as an indicator to understand the localization accuracy of our system.

5.2 Impact of Interval Filtering

To understand the impact of the measurement interval on localization performance, we evaluated RL-SAR over a partial implementation without interval filtering. The interval filtering process is described in §3.2.2. The filter is designed to meet the maximum gap requirement as explained in §2.2.2.

We applied the random trajectory strategy to 5 Line-of-Sight tags in the environment. We repeated scanning the tags until their localization confidence met the

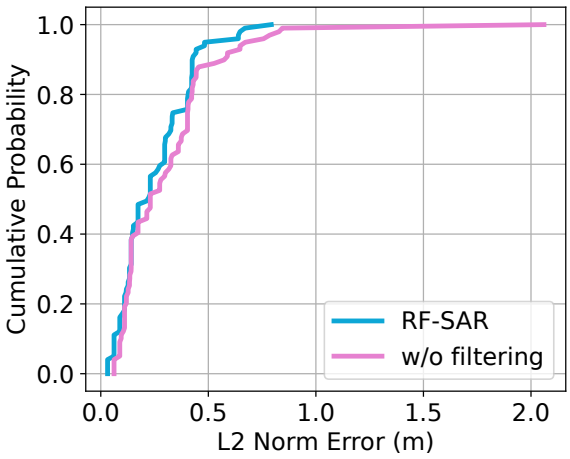

Figure 5-2: **Benefit of Interval Filtering.** This graph shows the CDF comparison of 3D localization error between the RL-SAR system with and without interval filtering. The system without interval filtering displays 3D localization error of more than 2m at high percentile, which might be caused by multiple localization candidates discussed in §2.2.2.

convergence criteria. Finally, we plotted the CDF of the localization accuracy for both implementations in Figure 5-2. The blue line shows the CDF of the results using RL-SAR, and the pink line displays the results using the partial implementation without interval filtering. We draw the following conclusions from the results:

- Both implementations of with and without interval filtering yield a similar median localization error of 0.223m and 0.230m, respectively.
- The implementation without an interval filter has a larger localization error for its 90^{th} percentile and above. The maximum localization error for this implementation is larger than 2m. The reason why the localization result converges to this poor estimate might be that the spatial aliasing effect creates a faulty localization candidate for the RF SAR hologram.

This experiment demonstrates the benefit of using interval filtering. It shows that this technique leads to better overall localization accuracy by preventing multiple localization candidates in the SAR computation.

Chapter 6

Results

6.1 Localization Accuracy

To understand the overall performance of RL-SAR, we evaluated the 3D localization accuracy of 15 tags in our experimental setup. These 15 tags were placed according to the profile defined in Table 4.1. The ground truths were carefully taken and measured relative to the origin. We repeatedly scanned the environment 10 times until the localization confidence converged for all the tags. Figure 6-1 plots the cumulative probability of the absolute error in X, Y, and Z directions. It indicates that the median X error is 0.07m, the median Y error is 0.07m, and the median Z error is 0.11m. Figure 6-2 combines the error in these three directions and plots the CDF of the total 3D localization error. It demonstrates that the median L2-norm error of RL-SAR is 0.225m.

We plot the change in localization error as more scans are performed and the total trajectory length increases. Figure 6-3 displays the localization error over total trajectory lengths, ranging from 6 to 27 meters. It reveals that RL-SAR has a mean error of 0.325m and a 90^{th} percentile of 0.763m when the scanning distance is 6 meters. As RL-SAR collects more measurements and the scanning distance increases, the mean and the 90^{th} percentile localization errors decrease to 0.248m and 0.333m, respectively. The trajectory length starts at 6 meters due to the initial trajectory scan τ_{init} that RL-SAR performs. For localization trials that converge before reaching 27

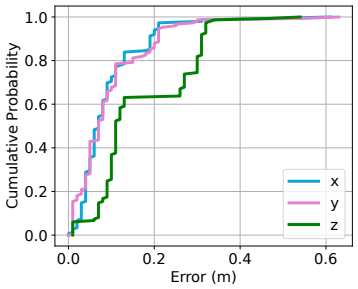

Figure 6-1: **CDF of Localization Error in XYZ axes.** This plot shows the cumulative probability of absolute error measurements in X, Y, and Z directions.

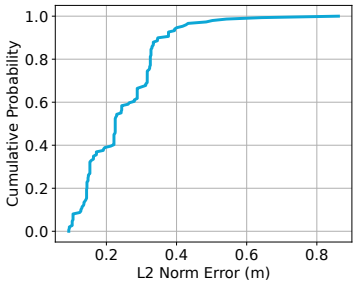

Figure 6-2: **CDF of Total 3D Localization Error.** This plot shows the cumulative probability of the total 3D localization error of RL-SAR.

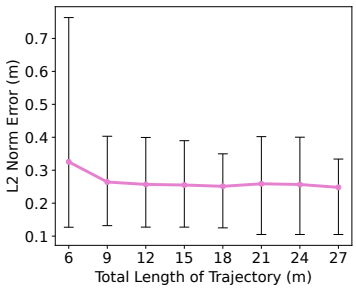

Figure 6-3: **L2-Norm Error over Trajectory Lengths.** This line plot demonstrates the mean L2-norm error at each scanning distance. The error bars mark the 10^{th} and 90^{th} percentiles.

meters, we extended the localization error at the time of convergence up to 27 meters in order to calculate the mean error across trials.

We compare the L2-norm localization error between the two naive strategies and the RL-SAR method. Figure 6-4 depicts the mean localization error for these three different strategies. We set a timeout of 20 steps, meaning that if the localization results don't converge within 20 steps, the system will utilize the last position estimation as the final tag observation. With the exhaustive scanning method, the mean 3D localization error is 0.456m and the 90^{th} percentile is 1.072m. The random trajectory strategy resulted in a mean error of 0.478m and a 90^{th} percentile of 1.195m. Finally, RL-SAR achieves a mean localization error of 0.244m and a 90^{th} percentile of 0.349m. Based on the results, we make the following remarks:

- RL-SAR exhibits a higher mean Z error (0.11m) than mean X (0.07m) and Y error (0.07m). This is due to the ceiling-mounted track system moving solely in a 2D plane, meaning the RF measurements lack diversity in the Z direction to increase the Z aperture.

- As shown in Figure 6-3, the localization error decreases as RL-SAR performs more scans. It shows that RL-based trajectory optimization algorithm continue to advise new scanning paths that improve the localization accuracy.

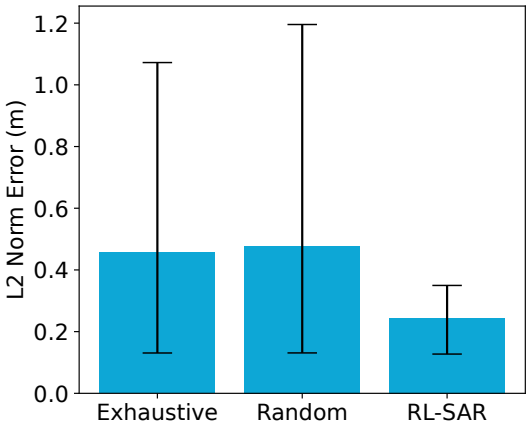

Figure 6-4: **Comparison of Localization Errors.** This bar plot shows the mean localization error across three different methods. The error bars indicate the 10^{th} and the 90^{th} percentiles of the error data.

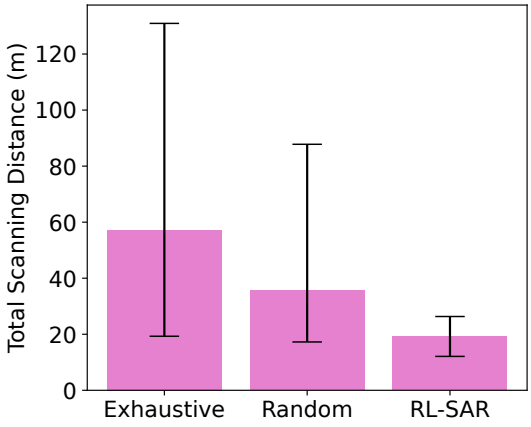

Figure 6-5: **Comparison of Total Scanning Distances.** This bar plot shows the total scanning distance across three different methods. The error bars indicate the 10^{th} and the 90^{th} percentiles of the distance data.

- Figure 6-4 shows that both the exhaustive and random trajectory strategies yield a higher mean error than RL-SAR. This could be due to the fact that these strategies do not optimize the path to obtain measurements that would enhance localization confidence. Therefore, the error tends to be larger, and in some trials, the localization results may not even converge before reaching the timeout steps.

Overall, the results indicate that RL-SAR has a mean localization error of 0.244m, and the RL-based trajectory optimization algorithm helps improve system accuracy with more scans. In addition, it has a 49% improvement in mean accuracy compared to the random trajectory strategy.

6.2 Total Scanning Distance

The primary goal of RL-SAR is to more efficiently locate multiple tags using the RL-based trajectory optimization algorithm. We compare our RL-based algorithm with two naive strategies: exhaustive search and random trajectory, as defined in §4.3. We implement these three methods in our experimental setup with 15 RFID tagged items

and repeat each strategy across 10 rounds. Each round concludes when all 15 tags converge or when they reach a predetermined timeout set to 20 steps. The results are plotted in Figure 6-5, which shows the mean total scanning distances of RL-SAR and their respective 10^{th} and 90^{th} percentiles for each strategy. We offer the following observations:

- The exhaustive search method yields the highest mean total scanning distance of 57.26m and a 90^{th} percentile of 130.93m. This is because it attempts to scan all spaces within the environment, leading to many inefficient long scans.
- The random trajectory strategy records the second highest mean total scanning distance of 35.59m and a 90^{th} percentile of 87.79m. In contrast, RL-SAR has a mean error of 19.14m and a 90^{th} percentile of 26.33m. It shows that the random trajectory strategy requires 86% longer scanning distance to achieve localization convergence on 15 RFID tagged items.

The results show that the baseline requires 86% longer trajectory than RL-SAR. It shows that the RL-based trajectory optimization algorithm improves the overall efficiency to find multiple RFID tags.

Chapter 7

Digital Twin Application

In this chapter, we show how RL-SAR can be leveraged to create digital twins of store-like environments. This application aims to create a digital map of items by merging the localization results from RL-SAR with a visual 3D reconstruction, using a depth camera. Such a digital twin would enable more efficient inventory tracking with intuitive visualization capabilities.

We built the application based on the Robot Operating System 2 (ROS2) framework [42]. In ROS2, data are communicated through messages. Each ROS node either publishes or subscribes to these messages for further processing. Real-time position feedback from RL-SAR and RF measurements are relayed back to the local server via a ZMQ socket. We created a ROS node to convert the ZMQ message into a ROS2 message, enabling integration with our application. Following each scan, the ROS node consolidates the RF measurements and computes SAR based on the method described in §3.2.3. The estimated location and confidence interval are then used to publish a visualization marker message, which can be used in the graphical user interface (GUI) integration.

We utilized the Intel Realsense D435i depth camera to record depth images of the environment. The camera, installed adjacent to the antenna, is directly connected to the local server through a USB 3.0 cable. We utilized the realsense-ros node [43] to receive real-time point cloud messages from the camera and developed a ROS node to aggregate multiple point cloud data as RL-SAR moves. This process creates a global

Figure 7-1: **Visualization of a Digital Twin.** RL-SAR creates a digital map of the environment and visualize the tags' locations on top of the 3D point cloud data of the environment. The screenshot on the left shows the visualization of this 3D map, and the blue circle defines the area that we primarily put the tags in. The images on the right shows the comparison between the actual setup and the digital visualization of the tags and their confidence intervals (green ellipses).

point cloud map that offers a detailed 3D reconstruction of the environment. Our ROS node then publishes this global point cloud for visualizing the digital twin.

To create a digital twin of the indoor space, we gathered localization result messages and global point cloud messages, and visualized them in a ROS2 GUI software named RVIZ2 [44]. The visualization is displayed in Figure 7-1. In RVIZ2, a 3D representation of the environment is shown, with green ellipses representing the estimated locations of our target tags. The length, width, and height of the ellipse denote the confidence intervals (c_x, c_y, c_z) of the specific tag. Furthermore, we developed a Universal Robot Description File (URDF) for the RL-SAR model. RVIZ2 visualizes this URDF file and synchronizes the movement between the actual hardware and the GUI. As RL-SAR continues scanning, the ellipses decrease in size as the localization results converge.

The digital twin application demonstrates the potential of a fine-grained RFID localization system in creating a real-time digital map of items within an indoor space.

Chapter 8

Conclusions

In this thesis, we introduced RL-SAR, an end-to-end, autonomous SAR-based RFID localization system featuring an RL trajectory optimization algorithm. We have implemented the system using an antenna that moves along a ceiling-mounted 2D track, collecting RF measurements from various locations. These measurements are processed to create a SAR hologram and estimate the location of each tag. The primary objective of RL-SAR is to identify the optimal locations for taking measurements to enhance localization confidence for multiple tags. We developed a RL algorithm based on Proximal Policy Optimization (PPO) to minimize the total distance that the antenna moves to obtain accurate localization results. Our evaluation results demonstrated that RL-SAR has a mean localization error of 0.244m, and that the naive strategies require 86% longer trajectory than RL-SAR. This shows the enhancement in localization efficiency for multiple RFID-tagged items. Furthermore, we showcased the potential for applying RL-SAR in real-world scenarios by implementing a digital twin application, creating a digital map of the environment using our system and a depth camera. As this research evolves, it would be interesting to explore how advancing the localizing algorithm with RF-visual sensor fusion (similar to [14]) and with more advanced RFID estimation techniques (similar to [2]) can further improve efficiency.

Bibliography

[1] C. Swedberg, "Rain rfid research report finds growth at 36 percent." [Online]. Available: https://www.rfidjournal.com/rain-rfid-research-report-finds-growth-at-36-percent

[2] Y. Ma, N. S. Selby, and F. M. Adib, "Minding the billions: Ultra-wideband localization for deployed rfid tags," *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, 2017.

[3] J. Wang and D. Katabi, "Dude, where's my card? rfid positioning that works with multipath and non-line of sight," *ACM SIGCOMM Computer Communication Review*, vol. 43, pp. 51–62, 08 2013.

[4] A. Motroni, P. Nepa, V. Magnago, A. Buffi, B. Tellini, D. Fontanelli, and D. Macii, "Sar-based indoor localization of uhf-rfid tags via mobile robot," *2018 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pp. 1–8, 2018.

[5] P. Tripicchio, M. Unetti, S. D'Avella, A. Buffi, A. Motroni, F. Bernardini, and P. Nepa, "A synthetic aperture uhf rfid localization method by phase unwrapping and hyperbolic intersection," *IEEE Transactions on Automation Science and Engineering*, vol. 19, pp. 933–945, 2021.

[6] C. Li, E. Tanghe, D. Plets, P. Suanet, J. Hoebeke, E. D. Poorter, and W. Joseph, "Reloc: Hybrid rssi- and phase-based relative uhf-rfid tag localization with cots devices," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, pp. 8613–8627, 2020.

[7] R. Zhao, D. Wang, Q. Zhang, H. Chen, and H. Xu, "Pec: Synthetic aperture rfid localization with aperture position error compensation," *2019 16th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pp. 1–9, 2019.

[8] Y. Jiang, Y. Ma, H. Liu, and Y. Zhang, "Rf-sml: A sar-based multi-granular and real-time localization method for rfid tags," *Electronics*, 2020.

[9] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus, "Rf-compass: Robot object manipulation using rfids," in *Proceedings of the 19th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '13. New

York, NY, USA: Association for Computing Machinery, 2013, p. 3–14. [Online]. Available: https://doi.org/10.1145/2500423.2500451

[10] L. Shangguan and K. Jamieson, "The design and implementation of a mobile rfid tag sorting robot," in *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 31–42. [Online]. Available: https://doi.org/10.1145/2906388.2906417

[11] R. Miesen, F. Kirsch, and M. Vossiek, "Uhf rfid localization based on synthetic apertures," *IEEE Transactions on Automation Science and Engineering*, vol. 10, pp. 807–815, 2013.

[12] A. Buffi, A. Motroni, P. Nepa, B. Tellini, and R. Cioni, "A sar-based measurement method for passive-tag positioning with a flying uhf-rfid reader," *IEEE Transactions on Instrumentation and Measurement*, vol. 68, pp. 845–853, 2019.

[13] A. Tzitzis, S. Megalou, S. Siachalou, E. Tsardoulias, T. V. Yioultsis, and A. G. Dimitriou, "3d localization of rfid tags with a single antenna by a moving robot and "phase relock"," *2019 IEEE International Conference on RFID Technology and Applications (RFID-TA)*, pp. 273–278, 2019.

[14] T. Boroushaki, I. Perper, M. Nachin, A. Rodriguez, and F. Adib, "Rfusion: Robotic grasping via rf-visual sensing and learning," in *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, 2021, pp. 192–205.

[15] F. Tlili, N. Hamdi, and A. Belghith, "Accurate 3d localization scheme based on active rfid tags for indoor environment," *2012 IEEE International Conference on RFID-Technologies and Applications (RFID-TA)*, pp. 378–382, 2012.

[16] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil, "Landmarc: Indoor location sensing using active rfid," *Wireless Networks*, vol. 10, pp. 701–710, 2004.

[17] N. B. Aldin, E. Erçelebi, and M. Aykaç, "An accurate indoor rssi localization algorithm based on active rfid system with reference tags," *Wireless Personal Communications*, vol. 97, pp. 3811–3829, 2017.

[18] J. Zhang, Y. Lyu, J. Patton, S. C. G. Periaswamy, and T. A. Roppel, "Bfvp: A probabilistic uhf rfid tag localization algorithm using bayesian filter and a variable power rfid model," *IEEE Transactions on Industrial Electronics*, vol. 65, pp. 8250–8259, 2018.

[19] A. Almaaitah, K. Ali, H. S. Hassanein, and M. Ibnkahla, "3d passive tag localization schemes for indoor rfid applications," *2010 IEEE International Conference on Communications*, pp. 1–5, 2010.

[20] Z. Luo, Q. Zhang, Y. Ma, M. Singh, and F. Adib, "3d backscatter localization for fine-grained robotics," in *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, 2019, pp. 765–782.

[21] P. V. Nikitin, R. Martinez, S. Ramamurthy, H. G. Leland, G. Spiess, and K. V. S. Rao, "Phase based spatial identification of uhf rfid tags," *2010 IEEE International Conference on RFID (IEEE RFID 2010)*, pp. 102–109, 2010.

[22] L. Yang, Y. Chen, X. Li, C. Xiao, M. Li, and Y. Liu, "Tagoram: real-time tracking of mobile rfid tags to high precision using cots devices," *Proceedings of the 20th annual international conference on Mobile computing and networking*, 2014.

[23] A. Buffi, P. Nepa, and F. Lombardini, "A phase-based technique for localization of uhf-rfid tags moving on a conveyor belt: Performance analysis and test-case measurements," *IEEE Sensors Journal*, vol. 15, pp. 387–396, 2015.

[24] Z. Liu, Z. Fu, T. Li, I. A. White, R. V. Penty, and M. J. Crisp, "A fast phase and rssi-based localization method using passive rid system with mobile platform," *2021 IEEE International Conference on RFID Technology and Applications (RFID-TA)*, pp. 48–51, 2021.

[25] A. Tzitzis, A. Filotheou, S. Siachalou, E. Tsardoulias, S. Megalou, A. Bletsas, K. L. Panayiotou, A. L. Symeonidis, T. V. Yioultsis, and A. G. Dimitriou, "Real-time 3d localization of rfid-tagged products by ground robots and drones with commercial off-the-shelf rfid equipment: Challenges and solutions," *2020 IEEE International Conference on RFID (RFID)*, pp. 1–8, 2020.

[26] T. Boroushaki, M. Lam, L. Dodds, A. Eid, and F. Adib, "Augmenting augmented reality with non-line-of-sight perception," in *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, 2023.

[27] T. Boroushaki, M. Lam, W. Chen, L. Dodds, A. Eid, and F. Adib, "Exploiting synergies between augmented reality and rfids for item localization and retrieval," in *IEEE International Conference on RFID (IEEE RFID 2023)*, 2023.

[28] L. Dodds, I. Perper, A. Eid, and F. Adib, "A handheld fine-grained rfid localization system with complex-controlled polarization," in *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, ser. ACM MobiCom '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: https://doi.org/10.1145/3570361.3592504

[29] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.

[30] J. Wang, D. Vasisht, and D. Katabi, "Rf-idraw: Virtual touch screen in the air using rf signals," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 235–246. [Online]. Available: https://doi.org/10.1145/2619239.2626330

[31] J. A. N. Company., "Thingmagic® m6e uhf rain rfid." [Online]. Available: https://www.jadaktech.com/product/thingmagic-m6e-uhf-rain-rfid/

[32] T. Dien Hoa, F. Ferrero, and L. Deneire, "Direction of arrival (doa) for 180 degree phase ambiguity rfid uhf reader: Limits and perspectives," 12 2017.

[33] J. A. N. Company., "Thingmagic® mercury api." [Online]. Available: https://www.jadaktech.com/product/thingmagic-mercury-api

[34] S. H. Doong, "A closed-form formula for GPS GDOP computation," *GPS Solutions*, vol. 13, no. 3, pp. 183–190, Jul. 2009. [Online]. Available: http://link.springer.com/10.1007/s10291-008-0111-2

[35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 07 2017.

[36] N. E. U. Inc., "Stepper motor – nema 23." [Online]. Available: https://us.nanotec.com/products/497-st5918-stepper-motor-nema-23

[37] T. G. Project, "What is grbl?" [Online]. Available: https://www.grbl.org/what-is-grbl

[38] "Mti mt-242025/trh/a (rhcp) outdoor rfid antenna (865-956 mhz)." [Online]. Available: https://www.atlasrfidstore.com/mti-mt-242025-trh-a-rhcp-outdoor-rfid-antenna-865-956-mhz/

[39] ZMQ, "Zeromq. an open-source universal messaging library." [Online]. Available: https://zeromq.org/

[40] J. Griffin and G. Durgin, "Complete link budgets for backscatter-radio and rfid systems," *Antennas and Propagation Magazine, IEEE*, vol. 51, pp. 11 – 25, 05 2009.

[41] "Stable-baselines3 docs - reliable reinforcement learning implementations." [Online]. Available: https://stable-baselines3.readthedocs.io/

[42] "Ros2 documentation." [Online]. Available: https://docs.ros.org/en/humble/index.html#

[43] "Intel(r) realsense(tm) ros wrapper for depth camera." [Online]. Available: https://github.com/IntelRealSense/realsense-ros

[44] "Ros 3d robot visualizer." [Online]. Available: https://github.com/ros2/rviz