Enhancing the Robustness of AI-Driven Robotic RFID Inventory Management Using Conformal Prediction

†Yongshuai Wu, †Jian Zhang, †Shaoen Wu, and ‡Shiwen Mao
†Department of Information Technology, Kennesaw State University, Marietta, GA 30060, USA
‡Department of Electrical and Computer Engineering, Auburn University, Auburn, AL 36849-5201, USA
Email: ywu26@students.kennesaw.edu, jianzhang@ieee.org, swu10@kennesaw.edu, smao@ieee.org

***Abstract*—In this work, we present a novel approach to enhance the robustness of autonomous robotic Radio Frequency Identification (RFID) inventory systems using Conformal Prediction (CP). Recent AI-driven approaches, especially deep-learning models, have made significant advances in performing inventory strategies and action planning. However, these models lack the capability to measure uncertainty during the prediction process, which can result in accumulated errors and lead to catastrophic failures. To address the above challenge, we propose a confidence-guaranteed policy using CP to ensure reliable predictions in RFID inventory tasks. Our method focuses on managing the uncertainty in sub-goal estimation for a trained model, ensuring that predictions can meet or exceed a user-specific confidence level. We conduct extensive experiments to assess the proposed method by regulating an existing model and evaluate its effectiveness in identifying uncertain predictions. The experimental results demonstrate the effectiveness of our approach in improving both the reliability and efficiency of RFID inventory tasks, ensuring consistent and trustworthy operation.**

***Index Terms*—Robustness, Conformal Prediction (CP), Radio-frequency identification (RFID), Inventory.**

I. Introduction

The Radio-frequency Identification (RFID) technology provides a low-cost and efficient solution for inventory management [1], [2]. It offers a touchless and non-line-of-sight (Non-LoS) inventory to enable robots to automatically scan all RFID-tagged items in a large space (e.g., warehouses, retail stores) by navigating to cover the interest area [3]. While previous work has successfully applied deep learning models and robotic systems to automate RFID inventory tasks [4], [5], they often face challenges in robustness to environmental changes. These discrepancies increase the uncertainty of the trained model and cumulatively lead to incorrect predictions, resulting in inefficiencies and failed inventory tasks. An example highlighted in [5] involves a robot colliding with a shelf and leading to a failed task, a typical failure resulting from accumulated uncertainties throughout the prediction process.

To address these challenges, we propose a confidence-guaranteed policy for reliable prediction in RFID inventory tasks, utilizing Conformal Prediction (CP) – a statistical method that quantifies and manages uncertainty in real-time [6], [7]. Previous studies across various domains, including natural language processing [8], indoor localization [9], robot navigation [10], [11], have demonstrated the effectiveness of CP-based confidence methods. These studies also employ uncertainty assessment metrics to evaluate model performance. Building on this, we develop a confidence-guaranteed RFID inventory policy that uses CP to detect uncertainty in sub-goal predictions and enables proactive responses. The sub-goal enables our method in various platforms without being affected by the hardware of the platforms. Our approach leverages a small set of calibration samples collected from the specific environment and task conditions to provide well-calibrated, confident sub-goals, ensuring the safe and effective execution of RFID inventory tasks. When uncertainties are detected, our method allows the robot to recover from uncertain predictions and subsequently request assistance from a human operator if repeated recovery attempts fail. Fig. 1 visually illustrates and compares our proposed method to existing methods. This calibration-based method supports real-time adjustments based on environmental uncertainties, user-defined confidence levels, and an efficient recovery process, helping address trust gaps and maintain operational safety [12]. Furthermore, we introduce several quantification metrics to evaluate the proposed method and the uncertain performance of the trained model. The major contributions of this work are summarised as follows:

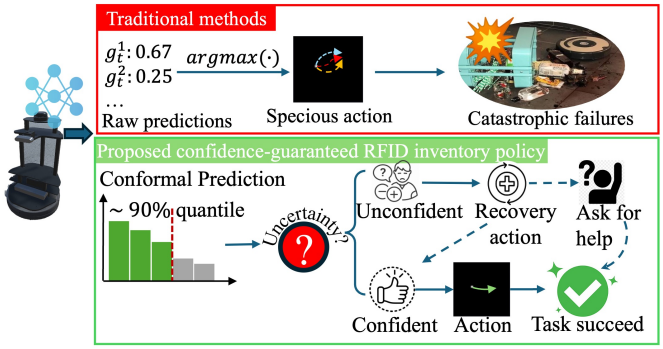

Fig. 1. Traditional AI-driven Methods vs. Proposed Confidence-Guaranteed RFID Inventory Policy: The figure's upper part illustrates traditional methods' limitations, where a lack of ability to detect uncertainty leads to specious actions, resulting in catastrophic failures. The bottom part shows the proposed confidence-guaranteed RFID inventory policy using Conformal Prediction, which efficiently detects prediction uncertainty and offers well-calibrated, confident predictions for conducting inventory tasks.

- This work proposes a confidence-guaranteed RFID in-

ventory policy by incorporating CP into a trained model, enhancing the model's robustness.
- We introduce the self-recovery action under model uncertainties to allow the robot to proactively collect additional information from environments or human users to gain a confident prediction. It enables the robot to smoothly and safely handle uncertain conditions during the deployment.
- Our method is based on calibration and task agnostic, making it more user-friendly in deployments. It creates a promising paradigm for robot deployment by collecting a small sample set from end users to ensure robots behave robustly in new environments.
- The proposed metrics evaluate an existing model's performance in terms of uncertainty, providing a solid tool for unveiling its internal robustness in specific scenarios that will help in future training.

The remainder of this paper is organized as follows: Section II outlines the problem statement and the objectives of this work. Section III provides a detailed explanation of the methods. Section IV-A describes the experimental setup, results, and analysis. Finally, the conclusions are presented in V.

II. Problem Statement

This work focuses on a robotic autonomous RFID inventory task $\mathcal{T}$, where a robot navigates in an unknown environment $\mathcal{E}$ to scan all RFID tags. The task will be broken down to visit a sequence of sub-goals and defined as $\mathcal{T} = \{g_0, \ldots, g_n\}$, with g_t is a sub-goal and n is the number of steps to complete the task. While the robot reaches a given sub-goal g_{t-1}, the next sub-goal will be predicted by a policy: $g_t = \pi(s_t)$. Here, $s_t = (o_t, \{g_i\}_{i=t-T-1}^{t-1})$ denotes the state at step t, which comprises T of historical sub-goals $\{g_i\}_{i=t-T-1}^{t-1}$ and the current observation o_t collected from environment $\mathcal{E}$. In this work, the robot perceives $\mathcal{E}$, and the collected raw data is processed by a pre-processing module introduced in [5] to produce the observation o_t. Then, a local planner will move the robot from g_{t-1} to g_t, using the built-in reader to scan RFID tags. Our ultimate goal in this project is to scan all RFID tags with an optimized route. For any given state s_t, it has and only has one ground truth sub-goal $\hat{g}_t$ that will navigate the robot in this optimized route. Several methods [5], [13] attempt to address this issue by learning the policy π and gaining good results in laboratory environments, but their performance degrades dramatically in deployments. Usually, the discrepancies between the deployment and training scenarios impose uncertainties to the trained policy, and the lack of ability to consider those uncertainties is the main reason for the performance degrades. To solve this robustness problem, we will achieve two goals in this work:

a) Goal 1. Develop a Confidence Guaranteed RFID Inventory Policy: To enhance the robustness of the RFID inventory task, we aim to develop a confidence guaranteed policy π_α that ensures the robot can always navigate to the ground truth sub-goal $\hat{g}_t$ as each step within an acceptable confidence margin. We formulate this goal as follows:

$$\mathbb{P}\big[\hat{g}_t \in \pi_\alpha(s_t)\big] \geq (1-\alpha), \quad \forall\ s_t \in \mathcal{T} \tag{1}$$

where $\alpha \in (0,1)$ is a user-specified error tolerance level, $\mathbb{P}$ denotes the probability. When satisfied, it guarantees that policy π_α can predict the ground truth $\hat{g}_t$ with $1-\alpha$ confidence.

b) Goal 2. Quantify uncertainty: While the $\mathcal{E}$ changes, it will result in a higher uncertainty for the policy. In this work, we aim to evaluate the uncertainty in our proposed π_α at any given step t by formulating the uncertainty function $\mathbb{U}$ as follows:

$$\mathbb{U}\Big(\pi_\alpha(s_t)\Big) = \begin{cases} 1, & \text{NO uncertainty detected} \\ 0, & \text{Uncertain} \end{cases} \tag{2}$$

this function indicates that the prediction result of π_α is considered "certain" when $\mathbb{U} = 1$. It enables the robot to be aware of its uncertainty in sub-goal prediction to avoid potential errors.

III. Proposed Method

This work introduces the confidence-guaranteed RFID inventory policy, π_a, a calibration-based method for regulating a well-trained model to new dynamic environments during deployment. This method enables the robot to measure its uncertainty in sub-goal prediction and ensure its action failure is consistently lower than the user-defined tolerance level, α, under different environments. In other words, it guarantees a success rate greater than the confidence level $1-\alpha$. We illustrate its architecture in Fig. 2: a well-trained model $\mathcal{M}$ will predict a goal probability map h_t at each step based on the state, a grid probability map h_t^g is then constructed from h_t, and then feed to the CP. The CP will generate a conformal prediction set to detect the uncertainty and predict the ground truth sub-goal $\hat{g}_t$.

A. Policy model

Inspired by our previous work [5], we deploy a method to allow the robot to learn a model $\mathcal{M} : s_t \rightarrow f(s_t) = h_t$

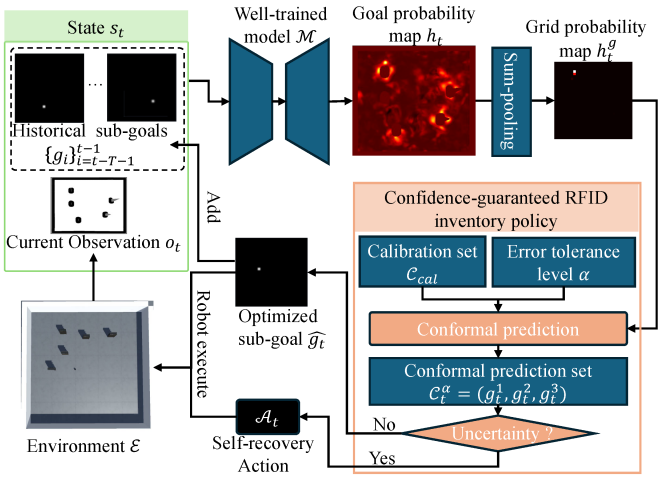

Fig. 2. The proposed confidence-guaranteed RFID inventory policy utilizes a small calibration set from the representative samples to detect sub-goal prediction uncertainties and act correspondingly. CP rigorously guarantees that $(1-\alpha)$ of the predicted sub-goals will be correct when without uncertainties detected. While uncertainties are detected, self-recovery actions will be offered to gain additional observations or ask for help from human users.

by giving "expert" samples $\{(s_0, g_0), \ldots, (s_m, g_m)|\mathcal{E}\}$. Here, $h_t \in \mathbb{R}^{H\times W}$ is the goal probability map with a width of W, and height of H, an example of h_t is given in Fig. 2. Each pixel in h_t corresponds to a location in the environment $\mathcal{E}$, with its value representing the model's empirical confidence. Considering the environmental noise and hardware defect, a sub-goal g_t is defined as a *navigational cycle zone* in h_t with noise level d_{err} as the radius. A robot arrives at g_t when it stops inside any place inside this zone. To settle this inherent noise, we apply a sum-pooling approach to process the h_t to generate a grid probability map $h_t^g : [0,1]^{H_g\times W_g}$, defined as:

$$h_t^g[x,y] = Softmax\Bigg(\sum_{i=x\cdot s_h}^{(x+1)\cdot s_h}\sum_{j=y\cdot s_w}^{(y+1)\cdot s_w} h_t[i,j]\Bigg) \quad (3)$$

where x, y is the coordinate of the h_t^g, s_h, s_w is the stride of the pooling window, and h_t^g is normalized with $Softmax$ function to value range from 0 to 1. The stride is chosen when the width or height of the h_t is divisible by the corresponding stride. It is a trade-off between navigation resolution and computational performance; a higher stride will have better computational performance but a lower navigation resolution. To achieve the best balance, usually, the stride should be in the range of $d_{err}/2$ to d_{err} based on empirical results. An example of this h_t^g is shown in the top right corner of Fig. 2. Thus, a sub-goal g_t will correspond to a specific position (x, y) in this grip map and can be defined as:

$$g_t = (x, y), \quad (4)$$

it coordinates with an empirical confidence value $h_t^g(x, y)$. Our previous experiment and initial investigations found that the empirical confidence may not be aligned with correctness: sometimes, the ground truth sub-goal $\hat{g}_t$ may not be the one with the highest value, especially when the test environment differs from the training environment.

B. Conformal Prediction based sub-goal prediction

To satisfy our goal (1), we employ CP to guarantee the confidence level of the model's prediction results. CP is a statistical method that converts any model's empirical conference value to a statistic rigorously guaranteed assurance [14], [15]. It utilizes a calibration set $\mathcal{C}_{cal}$ consisting of a small number of representative samples comprising the observations and correlated ground truth predictions for a given task in a specific environment. It regulates the model to ensure its output satisfies the given confidence level by calibrating the empirical conference value. More importantly, this rigorously guaranteed conformal prediction set $\mathcal{C}_t^\alpha$ can be produced at each step t to ensure it comprises the ground truth prediction with at least (1- α) confidence.

By deploying CP, we introduce a new robotic deployment paradigm. When a robot is deployed in a new environment for our RFID inventory task, end users first collect a small set of state samples $\mathcal{C}_{task} = \{s_0, \ldots, s_n\}$. Then, end users will select the ground truth sub-goal $\hat{g}_t$ for each s_t to form a calibration sample pair $(s_t, \hat{g}_t)$ and construct the $\mathcal{C}_{cal}$ as:

$$\mathcal{C}_{cal} = \{(s_0, \hat{g_0}), \ldots, (s_n, \hat{g_n})\}, s_t \in \mathcal{C}_{task} \quad (5)$$

This $\mathcal{C}_{cal}$ provides end users' specific expectations at the task context in this given environment. With the $\mathcal{C}_{cal}$, we will take several steps to calibrate any new prediction and ensure it is an optimized sub-goal with at least (1- α) probability.

Step 1. to calculate the conformal quantile. We start by calculating the empirical confidence value for each sample in $\mathcal{C}_{cal}$ and form the empirical confidence value set $\mathcal{C}_{ucv} = \{1 - h_t^g(\hat{g}_i)\}_{i=1}^n$. Here, $h_t^g(\hat{g}_i)$ is the empirical confidence value for the ground truth sub-goal $\hat{g}_i$ at i-th sample. Our grid probability map h_t^g at the position of $\hat{g}_i$ provides this empirical confidence value. With an user-specified error tolerance level α, we have $q = (n+1)(1-\alpha)/n$ to get the conformal quantile $\hat{q}$ of $\mathcal{C}_{cal}$ as:

$$\hat{q} = \mathbb{Q}(\mathcal{C}_{ucv}, \lceil q \rceil) \quad (6)$$

where $\mathbb{Q}$ is the quantile function, $\lceil\cdot\rceil$ is the ceil function.

Step2. Now with the $\hat{q}$, we can calculate the conformal prediction set $\mathcal{C}_t^\alpha$ for a new state $s_t \sim \mathcal{C}_{cal}$. Given the new state s_t, our model $\mathcal{M}$ produces the grid probability map h_t^g that offers candidate sub-goals and related empirical confidence value. Based on h_t^g, we can get the conformal prediction set $\mathcal{C}_t^\alpha$ as follows:

$$\mathcal{C}_t^\alpha = \{g_t | h_t^g(x, y) \leq \hat{q}\}_{x=1,y=1}^{H_g,W_g} \quad (7)$$

Where g_t is defined in (4). Therefore, the proposed conformal prediction based sub-goal prediction $\overline{\pi_\alpha}$ can be defined as:

$$\overline{\pi_\alpha}(s_t) = \mathcal{C}_t^\alpha, \forall s_t \sim C_{\text{cal}} \quad (8)$$

Based on the CP's statistical guarantee, we can satisfy our goal (1) with an additional constraint that the new state s_t should have a similar distribution with calibration set $\mathcal{C}_{cal}$.

C. Uncertainty Quantification & Self-recovery Actions

a) Uncertainty Quantification: Based on equation (8), the $\overline{\pi_\alpha}$ guarantees to produce a prediction set $\mathcal{C}_t^\alpha$ comprised of the ground truth subgoal within the given confidence level. From this $\mathcal{C}_t^\alpha$, we can quantify its prediction uncertainty at step t via analyzing the number of equivalent sub-goals. We formalize uncertainty quantification as follows:

$$\mathbb{U}_\alpha\big(\overline{\pi_\alpha}(s_t)\big) = \mathbb{I}\Big[\big|\mathbb{C}_l(\mathcal{C}_t^\alpha, d_{err})\big| = 1\Big], \quad (9)$$

where $|\cdot|$ denotes the cardinality of the set, $\mathbb{C}_l$ is a clustering function that measures the distances between candidate sub-goals in $\mathcal{C}_t^\alpha$ and groups those with distances less than d_{err} into an equivalent sub-goal, and the indicator function $\mathbb{I} = 1$ when only a single equivalent sub-goal is produced. The distance is the Euclidean distance measured in the units of the environment (such as meters). Thus, $\mathbb{U}_\alpha = 1$ indicates the prediction is certain because all elements in $\mathcal{C}_t^\alpha$ are in proximity within the noise level d_{err}; in other words, as previously stated, all these elements represent the equivalent ground truth sub-goal $\hat{g}_t$. Otherwise, we will have $\mathbb{U}_\alpha = 0$ indicates the policy is uncertain about the prediction result.

b) Self-recovery Action: By using the equations (8) and (9), we provide a way to detect task uncertainty, and the robot confidently performs the task when it provides a certain sub-goal. However, we still need a strategy to handle and respond to the uncertainties. To this end, we define and develop the self-recovery action $\mathcal{A}_t$ to allow the robot to recover from the uncertainties. Usually, $\mathcal{A}_t$ is a tasking-dependent strategy that needs the knowledge and context of a given task and environment. In our RFID inventory task, $\mathcal{A}_t$ will be deployed as a simple navigational strategy that allows the robot to move around the current position within a limited safe zone. This strategy empowers the robot to gather more environmental information, thereby making a certain prediction. We allow at most m_{rec} self-recovery attempts during the inventory task. If the model is still not able to predict a certain prediction, the robot will proactively ask the human user for help. Therefore, we complete our final **confidence-guaranteed RFID inventory policy** π_α with uncertainty quantification ability and self-recovery, and it is defined as follows:

$$g_t = \pi_\alpha(s_t) = \begin{cases} \mathbb{C}_l(\mathcal{C}_t^\alpha, d_{err}), & \text{when } \mathbb{U}_\alpha = 1 \\ \mathcal{A}_t, & \text{otherwise} \end{cases} \tag{10}$$

The π_α will output a specific sub-goal g_t while it is certain about the current prediction. Otherwise, it will produce the recovery action $\mathcal{A}_t$. Here, the probability of the g_t is an equivalent sub-goal of the ground truth $\hat{g}_t$ is at least $(1-\alpha)$.

D. Evaluation Metrics

Inspired by previous work [9], we proposed 3 metrics to quantify the uncertainty performance of a trained model $\mathcal{M}$. Given a evaluate set $C_{eval} = \{(s_0, \hat{g_0}), \ldots, (s_n, \hat{g_n})\}$, the Average Sub-goal Prediction set Size (ASP) is defined as:

$$ASP_\alpha(C_{eval}) = \frac{1}{n}\sum_{i=1}^{n} \left|\overline{\pi_\alpha}(s_i)\right|, \forall s_i \in C_{eval} \tag{11}$$

This metric represents the uncertain degree of the well-trained model $\mathcal{M}$ for the task; a large ASP_α denotes its lack of confidence. Then, we define the Uncertainty Detection Percentage (UDP) to describe the ratio of uncertain predictions:

$$UDP_\alpha(C_{eval}) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{U}_\alpha\left(\overline{\pi_\alpha}(s_i)\right), \forall s_i \in C_{eval} \tag{12}$$

UDP_α represents the overall confidence prediction ratio at the evaluation set; therefore, a UDP_α closer to 1 denotes that the model is more certain about the task. The Conformal Prediction Accuracy (CPA) represents the accuracy of conformal prediction, defined as:

$$CPA_\alpha = \frac{\sum_{i=1}^{n} \mathbb{I}\left[\mathbb{D}(\pi_\alpha(s_i), \hat{g}_i) < d_{err}\right]}{\sum_{i=1}^{n} \mathbb{I}\left[\mathbb{U}_\alpha = 1\right]}, \forall s_i \in C_{eval} \tag{13}$$

where $\mathbb{D}$ is the Euclidean distance function; as previously mentioned, a prediction $g_t = \pi_\alpha(s_t)$ is considered correct if it is within a distance d_{err} of the ground truth $\hat{g}_t$. CPA_α is defined as the rate of correct predictions among all certain predictions; thus, a CPA_α closer to 1 denotes that the proposed CP-based method effectively produces correct results in these certain predictions.

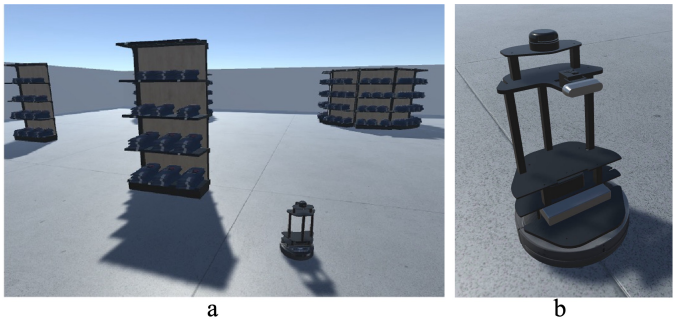

Fig. 3. An example of the virtual experimental environment: (a). It is a virtual indoor apparel store with a size of $5m \times 5m = 25m^2$, and 5 shelves are placed randomly in the room. (b). A zoom-in of the virtual robot that equipped a 2D Lidar and simulated RFID reader as main sensors.

IV. Experimental Study

A. Experiment Setup

In this section, we introduce the experimental environment and configurations for the proposed confidence-guaranteed RFID inventory policy π_α. We employed a well-trained model $\mathcal{M}$ similar to our prior work, which has sufficient capability to complete the autonomous RFID inventory task even in complex environments. For all experiments, we deployed a Unity3D-based digital-twin virtual environment proposed in [5] and adapted the new scenario in which the robot is asked to conduct an RFID inventory task in a novel environment that is not included in the training dataset. It was an indoor apparel store comprised of 5 shelves of 3 types, randomly placed in this $5m \times 5m = 25m^2$ room to create novel scenarios. We show an example of this virtual store in Fig. 3.a. As shown in Fig. 3.b, a virtual mobile robot was deployed in this store for inventory tasks. It is a virtual replica of an Interbotix X-Series LoCoBot Base robot, which is equipped with a 2D Lidar and a simulated RFID reader.

In this virtual environment, the navigational noise level $d_{err} = 0.3$ meters, which was calculated based on the navigation local planner precision and the localization accuracy of the virtual environment. The resolution of the goal probability map h_t was 0.05 meter per pixel. In our setup, all pixels in h_t were uniquely matched with the locations in the environment. To balance navigation resolution and computational performance, we chose a pooling window stride of $s_h = 8, s_w = 8$, resulting in a grid probability map h_t^g with the size of 32×32. Each grid in h_t^g was equal to an area with a length of 0.4 meters. Thus, a navigational cycle zone with $d_{err} = 0.3$ meters can be approximated as an area of 2×2 grids. For the confidence-guaranteed RFID inventory experiment, we set the maximum self-recovery attempts $m_{rec} = 3$. When the policy failed to yield a certain conformal prediction after the maximum attempts, an alert was sent to the user, and the user was asked to select the next sub-goal from the conformal prediction results. An episode was completed when the robot scanned all RFID tags in the room or reached the maximum step limit.

TABLE I
UNCERTAINTY QUANTIFICATION RESULTS

Error Tolerance Level α	ASP_α	UDP_α	CPA_α
0.10	444.59	0.186	0.967
0.15	348.99	0.262	0.956
0.20	274.87	0.315	0.948
0.25	207.92	0.361	0.930
0.30	111.07	0.434	0.899

B. Experiment Result and Analysis

a) Uncertainty Quantification: We first evaluated the uncertainty performance of the trained model $\mathcal{M}$ by analyzing the CP results of multiple task episodes. First, we conducted 15 episodes to collect and construct the calibration set $\mathcal{C}_{cal}$ with 368 samples that included state s_t and manually given ground truth sub-goal $\hat{g}_t$. Then, we used $\mathcal{C}_{cal}$ in our proposed confidence-guaranteed policy π_α to conduct other 135 episodes to collect 3287 samples, which formed the evaluation set C_{eval}. We also manually selected the true sub-goals for each sample in C_{eval}, then compared the results from π_α. Based on this C_{eval}, we calculated all the uncertainty quantification metrics defined in equations (11), (12), and (13). The experimental results are presented in Table. I.

The metric ASP_α offers an overall indicator of the uncertainties for $\mathcal{M}$ in step level. A higher value of ASP_α represents the model is more uncertain about its predictions, which is also affected by the user-defined error tolerance level α. Table. I tells a linearly decrease as α increases because when the users can tolerance more errors, the π_α utilizes a loose constraint that allows making more certain predictions with a higher risk of excluding the ground truth sub-goal in the predicted set, which is proofed by the decreasing CPA_α. UDP_α offers a tool to indicate the model $\mathcal{M}$'s uncertainty in the task level under the current environmental scenarios. A higher UDP_α indicates the model is more overall certain for the task. Table. I also shows the linearly increasing as the α increases, indicating that the model is more certain about the task under higher error tolerance. Experiment results show that by introducing the user-specific error tolerance α, the proposed π_α offers an effective method to allow users to adjust the model's confidence level to align with their expectations. With the statistical results of ASP_α and UDP_α, although our current model $\mathcal{M}$ can successfully execute the autonomous RFID inventory task at a very high scan percentage, without the regulation from our π_α its ability to handle uncertainty is limited that may accumulate leads to catastrophic failures. Additionally, the metric CPA_α shows the correctness ratio in certain predictions, and our π_α rigorously guarantees that certain predictions can offer a correctness rate that reaches or exceeds (1-α). The results in Table. I proved that our proposed π_α achieved this goal to empower robots to safely and confidently conduct RFID tasks in novel environments and can significantly reduce the risk of failures.

b) Insights of Uncertainty Identification: In this experiment, we dived into the details of uncertainty detection by comparing the behavior between the proposed confidence-guaranteed policy π_α and only the model $\mathcal{M}$ through conducting multiple RFID inventory tasks by them separately. We explored and discussed the behavior of both policies in uncertain prediction situations and analyzed how π_α regulated the behavior of the model $\mathcal{M}$ in episode-to-episode comparison to provide insights into the possible reasons that caused those behaviors through the intermediate result in the conformal prediction set $\mathcal{C}_t^\alpha$. Fig. 4.a demonstrates a situation when $\mathcal{M}$ predicts a sub-optimal sub-goal. The figure shows that the model can navigate from the initial position down along half of the circle shelf. However, it predicts a sub-goal that has already been navigated, represented with a yellow dot, instead of the optimal sub-goal, represented by a green dot. Fig. 4.b shows the conformal prediction result: the conformal prediction set in this position contains multiple candidate sub-goals, and we notice that this set also includes the optimal sub-goal. However, the π_α shows a higher uncertainty at this step, which outputs several candidate sub-goals over all racks. In contrast, Fig. 5 shows the behavior from π_α when it is certain about the prediction at other steps. Here, only one candidate sub-goal is also near our target scan rack. As we can see, π_α is more certain about navigating this type of shelf that only outputs a single sub-goal candidate, which indicates the $\mathcal{M}$ was better trained under this scenario.

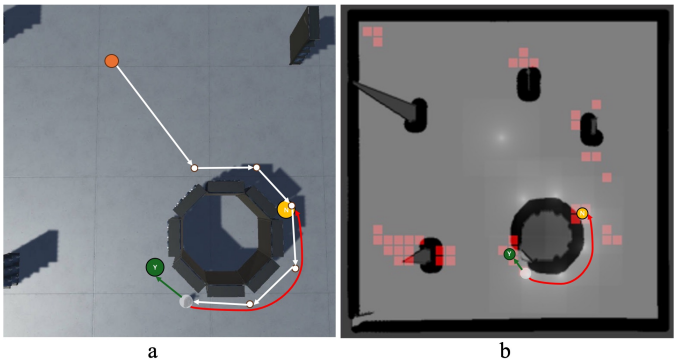

a b

Fig. 4. An example of insights into certainties detection. Both figures are in the same step $t = 13$; the green dot represents the optimal sub-goal, the yellow dot represents the predicted sub-goal, and the orange dot denotes the start point. (a) A close look at the wrong prediction made by only the $\mathcal{M}$ in the simulation environment with trajectory, (b) The map with CP prediction set $\mathcal{C}_t^\alpha$ from π_α, each candidate sub-goals are shown as red rectangles.

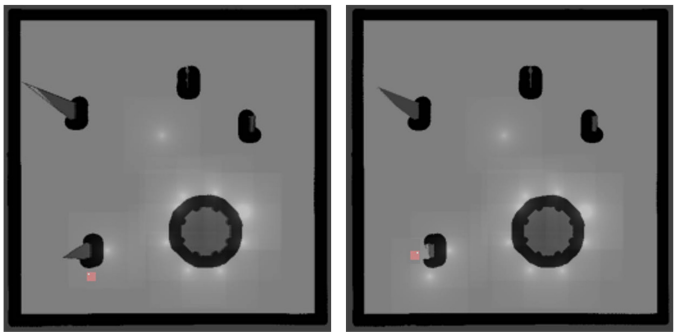

Fig. 5. Two sequential certain predictions provided by the proposed confidence-guaranteed policy π_α at steps $t = 16$ and $t = 17$.

TABLE II
UNCERTAINTY QUANTIFICATION RESULTS

Policy	Scanned Percentage	Average Travel Distance
$\mathcal{M}$	0.980	124.23
π_α	0.951	91.08

c) Confidence Guaranteed RFID inventory: This experiment evaluated the effectiveness of the proposed confidence-guaranteed policy π_a in the inventory tasks regarding RFID tag scanning. We ran $e = 25$ episodes of autonomous RFID inventory task that covers all possible shelf distributions for environment $\mathcal{E}$. Each environment was deployed with 288 RFID tags that were randomly placed on each rack. We evaluated its RFID inventory performance by comparing it with only the $\mathcal{M}$ and compared their results in Table II. We used the criteria in [5] to assess the robotic inventory task by the RFID tag scanned percentage and average travel distance for all episodes. An effective and efficient policy should offer a higher scanned percentage with short travel distances. Table II shows that the proposed π_α significantly reduces the average travel distance while maintaining the scanned percentage at the same level as the well-trained model $\mathcal{M}$, revealing the proposed policy's effectiveness. This experiment also unveiled the effectiveness of the proposed self-recovery action $\mathcal{A}_t$ in accomplishing the inventory task under uncertainties. Fig. 6 shows an episode of a completed confidence-guaranteed RFID inventory task while the robot adopted three recovery actions. In this episode, π_α demonstrated very high certainty when navigating a close shelf, but when it finished the shelf and was ready to transit, uncertainties arose. These uncertainties in the $\mathcal{M}$ might guide the robot randomly transited among racks, causing excessive time consumption; our proposed π_α identified them and proactively adopted $\mathcal{A}_t$. Then, the self-recovery action $\mathcal{A}_t$ enabled the robot to handle these uncertainties appropriately. Furthermore, this phenomenon indicates that the model $\mathcal{M}$ was poorly performed in transition scenarios and offers great clues to help in raining processes, such as providing more related training data.

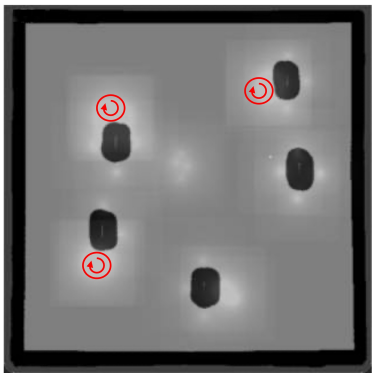

Fig. 6. An example of a confidence-guaranteed inventory task, the red cycles indicate the recovery actions raised and help accomplish the task under uncertainties detected.

V. CONCLUSION

This paper introduces a confidence-guaranteed RFID inventory management policy using Conformal Prediction (CP). The approach enhances the robustness of AI-driven robotic inventory systems by filtering out uncertain predictions and only allowing confident actions, improving task reliability and performance. Additionally, the proposed policy will proactively adopt self-recovery actions to handle the detected uncertainties and accomplish the inventory task with robust strategies. The evaluation metrics provide a systematic tool for assessing prediction uncertainty in any AI model. Extensive experiments confirm the method's effectiveness in managing uncertainties in dynamic environments. Future work will focus on integrating CP into training to improve performance in novel scenarios, such as for transitions between racks.

ACKNOWLEDGMENTS

The NSF partly supports this work under Grants CCSS-2245607 and CCSS-2245608.

REFERENCES

[1] D. Delen, B. C. Hardgrave, and R. Sharda, "Rfid for better supply-chain management through enhanced information visibility," *Production and operations management*, vol. 16, no. 5, pp. 613–624, 2007.
[2] J. Zhang, S. C. Periaswamy, S. Mao, and J. Patton, "Standards for passive uhf rfid," *GetMobile: Mobile Computing and Communications*, vol. 23, no. 3, pp. 10–15, 2020.
[3] J. Zhang, Y. Lyu, T. Roppel, J. Patton, and C. Senthilkumar, "Mobile robot for retail inventory using rfid," in *2016 IEEE international conference on Industrial technology (ICIT)*. IEEE, 2016, pp. 101–106.
[4] Z. Yu, J. Zhang, S. Mao, S. C. Periaswamy, and J. Patton, "Rirl: A recurrent imitation and reinforcement learning method for long-horizon robotic tasks," in *2022 IEEE 19th Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 2022, pp. 230–235.
[5] Y. Wu, J. Zhang, S. Wu, S. Mao, and Y. Wang, "Cmrm: A cross-modal reasoning model to enable zero-shot imitation learning for robotic rfid inventory in unstructured environments," in *GLOBECOM 2023-2023 IEEE Global Communications Conference*. IEEE, 2023, pp. 5354–5359.
[6] V. Vovk, "Conditional validity of inductive conformal predictors," in *Asian conference on machine learning*. PMLR, 2012, pp. 475–490.
[7] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," *arXiv preprint arXiv:2107.07511*, 2021.
[8] M. M. Campos, A. Farinhas, C. Zerva, M. A. Figueiredo, and A. F. Martins, "Conformal prediction for natural language processing: A survey," *arXiv preprint arXiv:2405.01976*, 2024.
[9] J. Ma, X. Wang, J. Zhang, S. Mao, S. C. Periaswamy, and J. Patton, "Navigating uncertainty: Ambiguity quantification in fingerprinting-based indoor localization," in *2024 IEEE Annual Congress on Artificial Intelligence of Things (AIoT)*. IEEE, 2024, pp. 123–128.
[10] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *IEEE Robotics and Automation Letters*, 2023.
[11] K. P. Kanhon, J. Zhang, and S. Mao, "Trustworthy hand signal communication between smart iot agents and humans," in *IEEE 10th World Forum on Internet of Thing*. Ottawa, Canada: IEEE, November 2024.
[12] A. Krausman, C. Neubauer, D. Forster, S. Lakhmani, A. L. Baker, S. M. Fitzhugh, G. Gremillion, J. L. Wright, J. S. Metcalfe, and K. E. Schaefer, "Trust measurement in human-autonomy teams: Development of a conceptual toolkit," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 11, no. 3, pp. 1–58, 2022.
[13] Z. Yu, J. Zhang, S. Mao, S. C. Periaswamy, and J. Patton, "Multi-state-space reasoning reinforcement learning for long-horizon rfid-based robotic searching and planning tasks," *Journal of Communications and Information Networks*, vol. 7, no. 3, pp. 239–251, 2022.
[14] M. Cauchois, S. Gupta, and J. C. Duchi, "Knowing what you know: valid and validated confidence sets in multiclass and multilabel prediction," *Journal of machine learning research*, vol. 22, no. 81, pp. 1–42, 2021.
[15] G. Shafer and V. Vovk, "A tutorial on conformal prediction." *Journal of Machine Learning Research*, vol. 9, no. 3, 2008.