Learning Visualization Policies of Augmented Reality for Human-Robot Collaboration

Anonymous Author(s)
Affiliation
Address
email

Abstract: In human-robot collaboration domains, augmented reality (AR) technologies have enabled people to visualize the state of robots. Current AR-based visualization policies are designed manually, which requires a lot of human efforts and domain knowledge. When too little information is visualized, human users find the AR interface not useful; when too much information is visualized, they find it difficult to process the visualized information. In this paper, we develop an intelligent AR agent that learns visualization policies (what to visualize, when, and how) from demonstrations. We created a Unity-based platform for simulating warehouse environments where human-robot teammates work on collaborative delivery tasks. We have collected a dataset that includes demonstrations of visualizing robots' current and planned behaviors. Our results from experiments with real human participants show that, compared with competitive baselines from the literature, our learned visualization strategies significantly increase the efficiency of human-robot teams in delivery tasks, while reducing the distraction level of human users.

Keywords: Augmented Reality, Multi-robot systems, Imitation Learning

1 Introduction

Robots are increasingly being used in industrial environments, such as manufacturing and warehouses, where the workspaces of people and robots are usually isolated from each other. **Human-robot collaboration (HRC)** opens up a plethora of opportunities where people and robots could complement their capabilities. Hence, robots are entering human-inhabited environments to work with humans. Humans and robots prefer different communication modalities. While humans are used to natural language and gestures in communication, robots exchange information in digital forms, such as text-based commands, resulting in a communication gap in HRC. Researchers have developed algorithms and system to bridge this communication gap using natural language [1, 2, 3] and vision [4, 5, 6]. **Augmented Reality (AR)** has been employed for HRC to provide an alternative communication mechanism with high bandwidth and low ambiguity [7].

AR is a technology that (1) combines real and virtual objects in a real environment; (2) runs interactively, and in real time; and (3) registers (aligns) real and virtual objects with each other [8]. Such composite views can enhance human perception of the real world, and also result in an interactive experience of the environment. Many researchers are working on uncovering the benefits of AR in HRC [9, 7]. For instance, researchers have developed AR systems that allow humans to visualize the state of the robots [10], as well as the robot intentions [11, 12]. Existing AR-based HRC systems employ a static visualization policy where the visualizations are always displayed to the user. The handcrafted static policies might suggest different visualization strategies, but those systems are not able to take important runtime factors into account (e.g., the number of robots, the current status of robots, the future intentions of robots, and the current status of human) to adapt the visualizations. The consequence is that those AR interfaces can alter people's attentional focus and result in *inattentional blindness* [13], if there is visual clutter caused by too many unwanted visualizations. Such concerns motivate our research on enabling our AR agent to learn a policy for dynamically selecting visualization actions given the state of the human-multi-robot system.

Submitted to the 6th Conference on Robot Learning (CoRL 2022). Do not distribute.

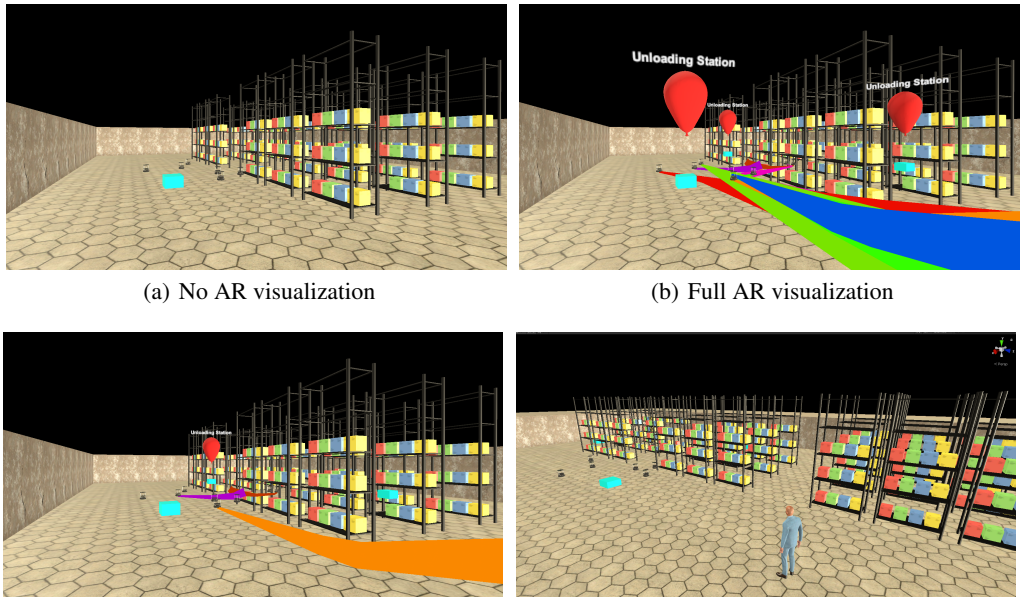

(a) No AR visualization (b) Full AR visualization

(c) Learned AR visualization (ours) (d) Third-person view of the environment

Figure 1: Different AR visualization strategies in a virtual warehouse environment, where a virtual human works with a team of mobile robots on collaborative delivery tasks. (a) No AR visualization, where the human worker does not know where some robots are, and what the robots plan to do. (b) Full AR visualization, where the human worker can be overwhelmed by the visual indicators. (c) Our learned AR visualization, where the AR agent uses a learned policy to dynamically determine a visualization strategy based on the current world state of both human and robots. (d) Third-person view of the human working alongside robots in the warehouse environment.

We design a framework called, ***Visualizations for Augmented Reality using Imitation Learning*** **(VARIL)**, for human-multi-robot collaboration in a shared environment. VARIL allows the human to track the status of a team of robots using an AR interface. In addition, VARIL learns the AR visualization policies using **Imitation Learning (IL)** [14] that dynamically selects AR visualization actions for human-multi-robot teams. To the best of our knowledge, this is the first work that employs learning to design visualization strategies of AR for human-multi-robot collaboration. In particular, we train decentralized visualization policies in a centralized fashion. While training, all the states are extracted from the agents in a centralized manner, and a new policy is generated by aggregating all the states, similar to learning in a single-agent environment. The new policy learned on the aggregated dataset is then deployed to the agents in a decentralized manner, where the agents query the actions from the joint policy, by using only the current agent's state information.

We empirically evaluated our VARIL framework in a warehouse domain, where both the human and the team of robots are engaged in a collaborative delivery task. Figure 1 shows our warehouse environment where twelve turtlebots collaborate with a human. [1] Both the human and robots are assigned nontransferable tasks, where the robots are tasked with delivering objects to different locations (drop stations) and the human is tasked with helping unload the objects from the robots waiting at drop stations. We collected a dataset of 6000 state-action pairs using the feedback provided from the "expert demonstrator". We have trained the AR agent in the simulated warehouse environment with this dataset using an iterative IL algorithm [15]. The learned AR visualization policies were compared with the static AR policy for human multi-robot teams from the literature [16], and the results suggest that the conflicts between the expert actions and the VARIL policies drastically decrease with the increase in the number of policy iterations. Additionally, we compared the total robot wait times for the intermediate policies of VARIL, and observed a significant decrease in the robot wait time between the initial and the learned AR visualization policies of VARIL.

[1] Our Unity-based simulated warehouse environment is hosted on Amazon Web Services at the following link, and will be open-sourced: `https://arroch-il.s3.us-east-2.amazonaws.com/index.html`

While we observed significant improvements in the human-robot collaboration efficiency in pure simulation, we decided to evaluate how VARIL influenced the user experience. We have evaluated VARIL with 25 participants where every participant operated a virtual human in the warehouse environment. Every participant was asked to fill out a survey questionnaire that was used for subjective evaluations. Based on the participants' responses, we observed that our learned visualization strategy significantly reduced the distraction caused by extraneous visualizations. In addition to the qualitative evaluation, we also compared the task completion times during the human study. The results show significant improvement in the efficiency of human-multi-robot teams in task completion as compared to two competitive baselines from the AR-for-HRC literature [17, 16].

2 Related Work

When humans and robots work in proximity, there is a need of communication to ensure the safety and efficiency of HRC. In this section, we first describe a few communication modalities used in current human-robot systems (non-AR and AR), and then discuss existing research that applies imitation learning methods to robotics domains.

In this subsection, we delineate a few communication modalities from the HRC literature. Researchers have used 2D interfaces in the past where the communication usually takes place using graphical user interfaces [18, 19, 20]. Also, natural language was used by a number of researchers for human-robot communication [21, 22]. Motion-based signaling of intent like gestures, including hand and facial have also been used for communication in human-robot teams [23, 24, 25]. Moreover, researchers studied how non-verbal social cues can improve task performance in human-robot teams [26]. Haptics, another communication modality, is largely based on physical interaction [27, 28, 29]. Even though all these modalities have their unique benefits and provide a medium for human-robot communication, the real world is 3D by nature, and these modalities are not strong in conveying spatial information, which can be very useful for HRC in shared spaces. Hence, despite those successes, AR has its unique advantages of visualizing and communicating spatial information with potentially less ambiguity and higher bandwidth [7].

Over a decade ago, researchers identified how AR can be used as a modality of communication when it comes to human-robot teams [7]. AR has the capability of overlaying spatial information onto the real world and helping a human counterpart visualize the robots' internal state and intended plans for effective HRC. Researchers have developed frameworks that enable human operators to visualize the motion-level intentions of unmanned aerial vehicles (UAVs) using AR [17, 30]. Zhu et al. augment the algorithms that control the autonomous robots to help understand the relationship between the robot's state and the environment around it [31]. In another line of research, Amor et al. used an AR interface to help humans visualize a robot arm's planned actions in the car assembly tasks [32]. Robot motion intent has been communicated and controlled using head-mounted displays through a mixed reality platform [33], which combines elements of augmented reality, virtual reality, and real-world objects.

Researchers have also developed a prototype system with a shared AR workspace that enables a shared perception, where a physical robot perceives the virtual elements in its own view and infer the utility of the human agent—the cost needed to perceive and interact in the workspace [34]. In another line, Chandan et al. have designed a system that enables the humans to visualize the state and intentions of a team of robots [16], where they further enabled human users to provide feedback to robot plans. More recently, researchers have designed different visual guidance in AR, where the humans can visualize the recommended actions from the robot (prescriptive guidance), and also visualize the state space information (descriptive guidance) to communicate why the robot recommends a particular action to the human [35]. All of the existing research employs static, predefined visualization policies. In contrast, we design a framework that learns a policy to select AR visualization actions (what to visualize, when, and how).

The integration of Imitation Learning (IL) with robotics has enabled the robots to learn a variety of tasks by observing expert demonstrations. One of the early applications of IL was to teach helicopters to fly very challenging maneuvers through providing expert demonstrations [36]. In another line of work, IL was used to enable a vehicle to mimic human driving behaviors in outdoor terrains [37, 38]. Later on, researchers used IL to learn from expert demonstrations in a simulated car driving domain [39]. A comprehensive study has been carried out to showcase the use of IL

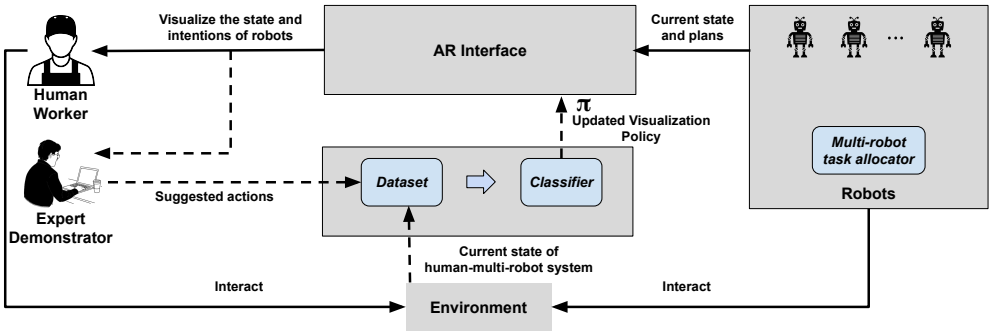

Figure 2: Overview of our VARIL framework that enables the human-multi-robot teams to collaborate in a shared environment. VARIL allows the human to track the status of a team of robots using an AR interface. Under the hood, VARIL supports learning an AR visualization policy using expert demonstrations, where the learned policy dynamically selects the actions for the AR agent.

in robotics [40]. Very recently, researchers have used AR interfaces for constrained Learning from Demonstration (LfD) to allow the agents to maintain, update, and adapt learned skills [41]. They have used learning techniques to enhance the skills of the robotic systems, whereas we train our AR agent to learn an adaptive visualization policy for human-multi-robot collaboration. Similar to the discussed existing research, we use IL to help our agent learn from expert demonstrations. Instead of learning to accomplish tasks directly, our VARIL framework leverages IL to improve the efficiency of AR-based HRC, while improving user experience by avoiding visual distractions from extraneous AR visualizations.

3 Framework

In this section, we describe our framework called ***Visualizations for Augmented Reality using Imitation Learning*** **(VARIL)**, and how it enables human-multi-robot teams to collaborate together in a shared environment. VARIL allows the human to track the status of a team of robots using an AR interface. Moreover, VARIL supports learning a decentralized AR visualization policy in a centralized fashion using expert demonstrations, where the AR agent uses the learned policy to dynamically select a visualization action based on the state of the human-multi-robot teams.[2] Figure 2 showcases the VARIL framework, and the interplay between the different components of VARIL.

Consider a team of robots working with a human worker in a shared space. The team of robots constantly share their state and plans using the AR interface, which is used by the human worker to track the status of the team of robots. The human worker simultaneously collaborates with the team of robots to complete the tasks. VARIL also consists of a human expert that gives demonstrations of AR visualizations at runtime, indicating what information should be visualized (or not) at specific times. Once a new policy is generated, the AR agent updates the visualizations to mimic the actions suggested by the expert demonstrator. There are two different human entities in VARIL, one for the role of the human worker, and the other is a human expert (Figure 2). Note that the expert is involved only during the policy learning phase, and once a satisfactory policy is learned, the expert demonstrator is no longer needed.

In the next subsection, we give a brief overview of how our human-multi-robot system works.

3.1 VARIL Algorithm

We design a human-multi-robot collaboration system using VARIL (Algorithm 1) that enables the humans to work with a team of robots while using an AR interface to track the status of the robot teammates. The inputs to the algorithm are as follows:

- $\mathcal{P}^t$: Task pool for the robots
- $\mathcal{P}^m$: Motion planner
- N: Number of robots
- $\mathcal{T}^a$: Multi-robot task allocator that assigns tasks at runtime

[2] Our "centralized training, decentralized execution" idea was inspired by recent multiagent reinforcement learning research [42], while it is applied to an IL domain in this work.

Algorithm 1 VARIL

Require: $\mathcal{P}^t, \mathcal{T}^a, \mathcal{P}^m, N, \hat{\pi}^{latest}$

1: Initialize $\hat{\pi}^{latest}$ to show all visualizations
2: Initialize empty lists: $\boldsymbol{\omega} \leftarrow \emptyset; \boldsymbol{\mathcal{I}} \leftarrow \emptyset$
3: Initialize $\mathcal{S}^{\mathcal{R}} \leftarrow \emptyset$, and $\mathcal{H}$ to $null$, where $\mathcal{S}^{\mathcal{R}}$ is the state of N robots, and $\mathcal{H}$ is the state of the human worker
4: $\boldsymbol{T} = \mathcal{T}^a(\mathcal{P}^t, N)$, where $t \in \boldsymbol{T}$ is a task assigned to one robot, and $|\boldsymbol{T}| = N$ ▷ More details on $\mathcal{T}^a$ in supplementary materials.
5: **while** $\mathcal{P}^t$ is **not** empty **do** ▷ Tasks still exist in task pool
6: **for** $i \in [0, 1, \cdots, N\text{-}1]$ **do**
7: **if** t_i is **not** complete **then**
8: Obtain i^{th} robot's configuration: $\boldsymbol{\omega}_i \leftarrow \theta(i)$
9: $\boldsymbol{\mathcal{I}}_i \leftarrow \mathcal{P}^m(\boldsymbol{\omega}_i, t_i)$
10: The ith robot follows $\boldsymbol{\mathcal{I}}_i$ using a controller
11: **else**
12: Update t_i using $\mathcal{T}^a$
13: **end if**
14: **end for**
15: $\lambda \leftarrow \mathcal{V}(\boldsymbol{\omega}, \boldsymbol{\mathcal{I}}, \hat{\pi}^{latest})$ ▷ More details on $\mathcal{V}$ in supplementary materials.
16: $\mathcal{S}^{\mathcal{R}} \leftarrow$ obtain all robots current state
17: $\mathcal{H} \leftarrow$ obtain current human worker state
18: $\hat{\pi}^{latest} \leftarrow$**PolicyUp**$(\mathcal{S}^{\mathcal{R}}, \mathcal{H})$ ▷ Update visualization policy
19: **end while**

Now, we describe in detail how our VARIL algorithm works. In Line 1, VARIL initializes $\hat{\pi}^{latest}$ to **enable all the visualizations**. Here visualizations represent all the 3D objects that are rendered in the AR interface to visualize the status of the team of robots. The intuition behind enabling all the visualizations is to enable the expert to easily identify different types of available visualizations. [3] The expert demonstrator can provide feedback on which visualizations are helpful, and which ones need to be turned on/off based on the observed state of the human-multi-robot system.

In Line 4, VARIL uses the multi-robot task allocator ($\mathcal{T}^a$) to allocate tasks (T) for N robots, where T_r is the task of the rth robot. Then, VARIL enters the main while-loop in Line 5 that runs until all the tasks from the task pool are completed. Inside the while-loop, VARIL enters a for-loop that runs for N times, once for each robot. At every iteration i of the for-loop, VARIL first checks if the robot has completed the current task or not (Line 7). If the task is not completed, VARIL first obtains the current robot configuration ($\boldsymbol{\omega}_i$). Then, VARIL uses $\mathcal{P}^m$ to generate the intended trajectory of the ith robot, $\boldsymbol{\mathcal{I}}_i$. Once, the current task is complete, the robot obtains its new task using $\mathcal{T}^a$.

After the for-loop, in Line 15, the set of robot configurations, robots' intended trajectories, and the latest visualization policy are passed to the AR agent that renders the visualizations using the AR interface. In Lines 16 and 17, VARIL obtains the current state of all robots ($\mathcal{S}^{\mathcal{R}}$), and the current state of the human worker ($\mathcal{H}$). Finally, VARIL calls the PolicyUp function that updates the global policy $\hat{\pi}^{latest}$. This enables VARIL's AR agent to update the visualizations based on $\hat{\pi}^{latest}$ in the next iteration. In the next subsection, we briefly explain how the PolicyUp function updates the visualization policy using IL based on the state of the human-multi-robot system and the expert demonstrations.

3.2 AR Policy Learning: PolicyUp

PolicyUp collects a dataset from the experiences of all the agents at each iteration, to train a new policy that best mimics the expert on the whole dataset. PolicyUp algorithm is an adapted version of **DAgger** algorithm [15], where PolicyUp updates the visualization policy using imitation learning based on the state of the human-multi-robot system, and the expert demonstrations. More details of the PolicyUp algorithm are provided in the supplementary materials.

[3] If we had turned off the AR visualizations, there would be the issue that the expert human could not locate the robots obscured by other objects.

4 Experiments

We conducted experiments in a simulated warehouse environment that we developed using Unity. With the experiments, we aim to evaluate **two hypotheses**: I) VARIL improves the overall efficiency in human-multi-robot team task completion, determined by the slowest agent of a team, in comparison to other AR-based methods from the literature that employ static visualization policies; and II) VARIL provides better user experience for humans in human-multi-robot collaboration tasks compared to competitive AR-based methods from the literature. We have evaluated both Hypothesis-I and Hypothesis-II with human participants collaborating with multiple robots in the simulated warehouse environment.

Participants: Twenty-five participants of ages 20-30, participated in the experiment, including six females and nineteen males. Every participant was given a $10 gift card for participating in the experiments. Each participant conducted three trials using three methods, including VARIL, that were randomly ordered for each participant. Before starting the actual trials, each participant completed a short dummy trial that was scaled down to a mini warehouse that includes only one robot. The role of the dummy trial was to acquaint the participants with the environment and to teach the participant to control the virtual human. The experiments were approved by the Institutional Review Board (IRB) of our institution.

4.1 Experiment Setup

We deployed the simulated warehouse environment to a web server to facilitate the online experiment using the WebGL build of the Unity project. In our experiments, the human participants were able to teleoperate the virtual human around the environment through the input of the W, A, S, and D keys along with using the mouse to rotate the virtual human's field of view. Each robot is assigned three boxes for each trial and once the last box has been dropped off by the robot, it will begin navigating to its starting position. The trial is completed once all twelve robots have delivered their three boxes and have reached their starting position.

Data Collection and Policy Learning: We collected a dataset of 6000 demonstrations provided by a human expert at runtime during which a human worker was working with a team of robots to collaborate on a delivery task. The human expert is one of the authors of this paper, who knows this domain well. The human worker and the human expert both were provided with an AR interface to track the robots. The human worker used the AR interface to collaborate with the robots, whereas the role of the human expert was to provide feedback on the visualizations seen using the AR interface.

The human expert was provided with a simple interface as shown in Figure 3(a) to input the feedback (which visualizations the human expert thinks should be to enabled or disabled) on the visualizations for the team of robots, as well as the drop station balloons. The interface provides a checkbox for each robot's trajectory, live location, and transparent avatar as well as a checkbox for each drop station balloon. The expert feedback is obtained in real time and the environment is not paused. The expert does not have access to the internal state representation, and the feedback is completely based on the human expert's observations on the environment. Along with the state of the human worker and the team of robots, the feedback (actions) of the human was combined to collect the dataset. A snapshot of the state-action pairs was extracted every four seconds, and after every 25 state-action pairs, the dataset was aggregated with the dataset from the previous iteration. Using VARIL, we ran a total of two hundred forty iterations, creating the entire dataset of the state-action pairs. At every iteration, we train the AR agent with a multi-class SVM using the all-pairs reduction to binary classification [43].

4.2 Baselines

For baselines, we have replicated the visualizations of two different systems from the literature, called ARROCH [16], and CRMIAR [17] [4] to make comparisons of their visualization strategies with the learned visualization from VARIL. ARROCH was designed for human-multi-robot systems, where a human can use an AR device to track the status of the robots. Our implementation

[4] We create the acronym of CRMIAR that indicates "Communicating Robot Motion Intent with Augmented Reality" for the simplicity of referring to the method.

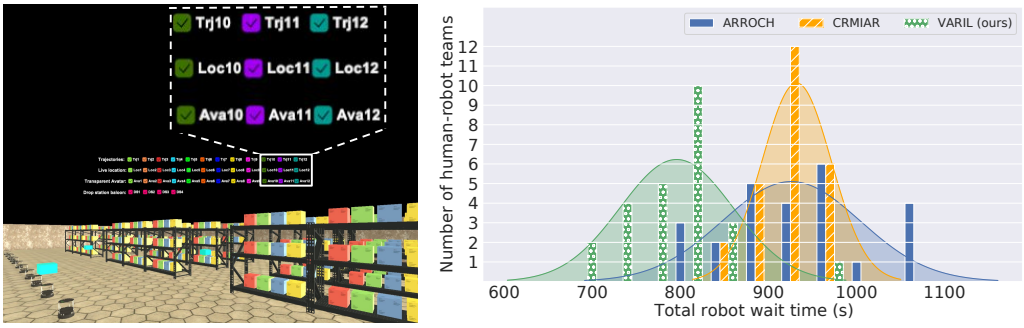

(a) (b)

Figure 3: (a) Interface used by the expert to give demonstrations. (b) Comparison of VARIL (our method) with two other methods from the literature to show the distribution of total waiting time (second) of all robots.

of the ARROCH system has the following visualizations: the trajectories to show the motion intentions, robot avatar to show the live location, and a transparent robot avatar to show the direction of the motion of the robot. On the other hand, CRMIAR was a single robot visualization system, but we have scaled it to enable tracking the team of robots in our warehouse environment. In our implementation of the CRMIAR system, where we have combined the NavPoints, and Utilties designs of CRMIAR to provide the following visualizations: the trajectories of robots to show their planned motion intentions, a map to show the relative location of the robots with respect to the human, and arrows to point to the robots that are not in the field of view. We use both ARROCH and CRMIAR in our experiments to evaluate the efficacy of VARIL. In the supplementary materials, we provide the figures of the both ARROCH and CRMIAR baseline implementations.

4.3 Results

Objective Results: Figure 3(b) shows the histogram, where the x-axis represents the total wait time of all the robots in seconds, and the y-axis represents the number of human-robot teams with the corresponding robot wait times. From the figure, it can be observed that for most human-robot teams, VARIL had the shortest robot wait times as compared to ARROCH and CRMIAR. Also, we plotted the Gaussian curves to show the distribution of the data points for all the methods. We also analyzed the statistical significance, in every trial, we sum up the task completion time of all the robots. We found that VARIL performed **significantly better** than both ARROCH and CRMIAR, where $0.01 < p{-}value < 0.05$. Additionally, by comparing the total robot wait times, we observed that the wait times for robots in VARIL were significantly shorter than the baselines. All of these results support Hypothesis-I that states VARIL improves the human-robot team's task completion efficiency.

Subjective Results: At the end of each trial, the participants were asked to fill out a survey form indicating their qualitative opinion over the six different questions to validate Hypothesis-II. The significant improvements observed from the study support Hypothesis-II on user experience that the visualizations in VARIL provided a better user experience. Collectively, the results convey that the VARIL visualizations help the participant better keep track of the robot status, proved useful towards task completion, and were less distracting, while the participants enjoyed using the system. The figure and the descriptive subjective analysis are provided in the supplementary materials.

Policy Evaluation: VARIL is the first work on Imitation Learning-based visualization policy learning for AR, and hence we do not have learning baselines to make comparisons for the generated policy. To overcome this shortcoming, we have evaluated our learned policy to find the discrepancy between the suggested expert actions as well as the actions suggested by both the ARROCH static visualization policy, and the intermediate policies of VARIL. Figure 4(a) shows the number of policy iterations on the x-axis, and the number of discrepancies on the y-axis. The discrepancy is a measure of the deviation of the visualization policies from the expert's policy. From the figure,

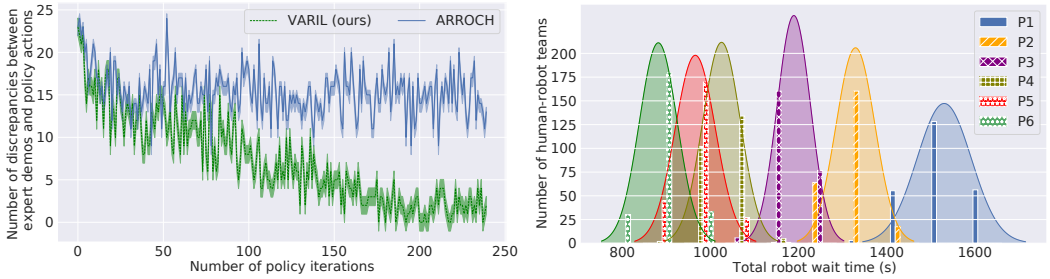

(a) Discrepancies: expert vs. policy

(b) Wait time (histogram)

Figure 4: (a) The number of discrepancies between the expert demonstrator's suggested actions with both the visualization policies of ARROCH (baseline) and VARIL with respect to the number of iterations. Each instance of a human expert disabling a visualization is considered a **discrepancy**; (b) A histogram to show the distribution of total waiting time (second) of all robots in simulation with respect to six intermediate policies of VARIL.

we can observe that the early policies of VARIL had more discrepancy, and this is because at the beginning the policy's poor quality makes it less effective in selecting visualization actions. With more interaction with the expert, the policy becomes mature, and tries to mimic the actions suggested by the expert. It can be observed that the discrepancies drastically decrease with the increase in the number of policy iterations. On the other hand, the increase of policy iteration has no effect on the number of discrepancies between ARROCH and expert actions, because ARROCH employs a static visualization policy. The final policy was used in the experiments to evaluate the learned visualization policy against the static policies from the baselines.

Furthermore, we evaluated Hypothesis-I in simulation to investigate how the intermediate policies of VARIL affect the human-robot team's task completion time. Figure 4(b) shows the histogram of the total robot wait times, for six different policies (P1 - P6), where every policy was extracted after 50 policy iterations. A total of 250 trials were run for every policy, and we observed improvements in the robot wait times denoting the improvement of efficiency in human-robot teams.

5 Conclusion and Future Work

In this paper, we present our framework, VARIL (*Visualizations for Augmented Reality using Imitation Learning*), that for the first time introduces a learning-based Augmented Reality (AR) visualization strategy for human-multi-robot collaboration. We have designed a system that learns a decentralized visualization policy using imitation learning in a centralized manner, where the AR agent dynamically selects a visualization action based on the state of the human-multi-robot system. We compared our framework with competitive baselines from the literature, and the results suggest that VARIL significantly increases the efficiency of human-robot collaboration, and reduces the distractions caused by extraneous information in AR.

Although there were human participants involved in this research, human-robot collaboration occurred only in simulated warehouses. Bringing VARIL to the real world will require extensive engineering efforts, and conducting experiments in a real warehouse is particularly challenging. To realize this system in the real-world, there are many difficulties, for example, the human worker needs to walk between drop stations, where each navigation action might take a long time. Adding to this complexity, the robots need to collaborate in realtime to complete the tasks where a single robot getting stuck to a dynamic obstacle can terminate the trial prematurely. To avoid such issues, we designed an open source simulator that can simulate a warehouse environment including a team of robots. Even though such a simulator does not overcome all the challenges for human-robot collaboration, we expect that this will enable AR-for-HRC researchers to focus on the crucial aspects of algorithm design rather than the tedious software development work. In the future, we plan to construct a built-in-lab mock warehouse environment for conducting real-world experiments.

References

[1] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. Teller, and N. Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.

[2] J. Y. Chai, L. She, R. Fang, S. Ottarson, C. Littley, C. Liu, and K. Hanson. Collaborative effort towards common ground in situated human-robot dialogue. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 33–40. ACM, 2014.

[3] C. Matuszek, E. Herbst, L. Zettlemoyer, and D. Fox. Learning to parse natural language commands to a robot control system. In *Experimental Robotics*, pages 403–415. Springer, 2013.

[4] S. Waldherr, R. Romero, and S. Thrun. A gesture based interface for human-robot interaction. *Autonomous Robots*, 9(2):151–173, Sep 2000.

[5] K. Nickel and R. Stiefelhagen. Visual recognition of pointing gestures for human–robot interaction. *Image and vision computing*, 25(12):1875–1884, 2007.

[6] H.-D. Yang, A.-Y. Park, and S.-W. Lee. Gesture spotting and recognition for human–robot interaction. *IEEE Transactions on robotics*, 23(2):256–270, 2007.

[7] S. A. Green, M. Billinghurst, X. Chen, and J. G. Chase. Human-robot collaboration: A literature review and augmented reality approach in design. *International journal of advanced robotic systems*, 5(1):1, 2008.

[8] R. Azuma, Y. Baillot, R. Behringer, S. Feiner, S. Julier, and B. MacIntyre. Recent advances in augmented reality. *IEEE computer graphics and applications*, 21(6):34–47, 2001.

[9] Z. Makhataeva and H. A. Varol. Augmented reality for robotics: a review. *Robotics*, 9(2):21, 2020.

[10] F. Muhammad, A. Hassan, A. Cleaver, and J. Sinapov. Creating a shared reality with robots. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction*, 2019.

[11] R. T. Chadalavada, H. Andreasson, R. Krug, and A. J. Lilienthal. That's on my mind! robot to human intention communication through on-board projection on shared floor space. In *2015 European Conference on Mobile Robots (ECMR)*. IEEE, 2015.

[12] A. Watanabe, T. Ikeda, Y. Morales, K. Shinozawa, T. Miyashita, and N. Hagita. Communicating robotic navigational intentions. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.

[13] B. J. Dixon, M. J. Daly, H. Chan, A. D. Vescan, I. J. Witterick, and J. C. Irish. Surgeons blinded by enhanced navigation: the effect of augmented reality on attention. *Surgical endoscopy*, 27 (2):454–461, 2013.

[14] S. Schaal. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3 (6):233–242, 1999.

[15] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.

[16] K. Chandan, V. Kudalkar, X. Li, and S. Zhang. ARROCH: Augmented Reality for Robots Collaborating with a Human. *In 2021 International Conference on Robotics and Automation (ICRA)*, 2021.

[17] M. Walker, H. Hedayati, J. Lee, and D. Szafir. Communicating robot motion intent with augmented reality. In *Proceedings of the International Conference on Human-Robot Interaction*, 2018.

[18] M. Baker, R. Casey, B. Keyes, and H. A. Yanco. Improved interfaces for human-robot interaction in urban search and rescue. In *2004 IEEE International Conference on Systems, Man and Cybernetics*, volume 3, pages 2960–2965, 2004.

[19] M. W. Kadous, R. K.-M. Sheh, and C. Sammut. Effective user interface design for rescue robotics. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 250–257. ACM, 2006.

[20] H. K. Keskinpala, J. A. Adams, and K. Kawamura. Pda-based human-robotic interface. In *2003 IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3931–3936, 2003.

[21] S. Amiri, S. Bajracharya, C. Goktolga, J. Thomason, and S. Zhang. Augmenting knowledge through statistical, goal-oriented human-robot dialog. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[22] J. Thomason, S. Zhang, R. J. Mooney, and P. Stone. Learning to interpret natural language commands through human-robot dialog. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[23] C. Breazeal. Emotion and sociable humanoid robots. *International journal of human-computer studies*, 59(1-2):119–155, 2003.

[24] K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, A. Sisbot, R. Alami, and T. Siméon. How may i serve you?: a robot companion approaching a seated person in a helping context. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 172–179. ACM, 2006.

[25] B. Robins, P. Dickerson, P. Stribling, and K. Dautenhahn. Robot-mediated joint attention in children with autism: A case study in robot-human interaction. *Interaction studies*, 5(2):161–198, 2004.

[26] C. Breazeal, C. D. Kidd, A. L. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *2005 IEEE/RSJ international conference on intelligent robots and systems*, 2005.

[27] T. Miyashita, T. Tajika, H. Ishiguro, K. Kogure, and N. Hagita. Haptic communication between humans and robots. In *Robotics Research*, pages 525–536. Springer, 2007.

[28] R. J. Stone. Haptic feedback: A brief history from telepresence to virtual reality. In *International Workshop on Haptic Human-Computer Interaction*, pages 1–16. Springer, 2000.

[29] J. Dumora, F. Geffard, C. Bidard, T. Brouillet, and P. Fraisse. Experimental study on haptic communication of a human in a shared human-robot collaborative task. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5137–5144. IEEE, 2012.

[30] H. Hedayati, M. Walker, and D. Szafir. Improving collocated robot teleoperation with augmented reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 78–86, 2018.

[31] D. Zhu and M. Veloso. Virtually adapted reality and algorithm visualization for autonomous robots. In *Robot World Cup*, pages 452–464. Springer, 2016.

[32] H. B. Amor, R. K. Ganesan, Y. Rathore, and H. Ross. Intention projection for human-robot collaboration with mixed reality cues. In *International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[33] E. Rosen, D. Whitney, E. Phillips, G. Chien, J. Tompkin, G. Konidaris, and S. Tellex. Communicating and controlling robot arm motion intent through mixed-reality head-mounted displays. *The International Journal of Robotics Research*, 38(12-13):1513–1526, 2019.

[34] S. Qiu, H. Liu, Z. Zhang, Y. Zhu, and S.-C. Zhu. Human-robot interaction in a shared augmented reality workspace. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11413–11418. IEEE, 2020.

[35] A. Tabrez, M. B. Luebbers, and B. Hayes. Descriptive and prescriptive visual guidance to improve shared situational awareness in human-robot teaming. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 1256–1264, 2022.

[36] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.

[37] J. Bagnell, J. Chestnutt, D. Bradley, and N. Ratliff. Boosting structured prediction for imitation learning. *Advances in Neural Information Processing Systems*, 19, 2006.

[38] D. Silver, J. Bagnell, and A. Stentz. High performance outdoor navigation from overhead data using imitation learning. *Robotics: Science and Systems IV, Zurich, Switzerland*, 1, 2008.

[39] J. Zhang and K. Cho. Query-efficient imitation learning for end-to-end autonomous driving. *arXiv preprint arXiv:1605.06450*, 2016.

[40] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard. Recent advances in robot learning from demonstration. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:297–330, 2020.

[41] M. Luebbers, C. Brooks, C. Mueller, D. Szafir, and B. Hayes. Arc-lfd: Using augmented reality for interactive long-term robot skill maintenance via constrained learning from demonstration. *In 2021 International Conference on Robotics and Automation (ICRA)*, 2021.

[42] T. Rashid, M. Samvelyan, C. Schroeder, G. Farquhar, J. Foerster, and S. Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 4295–4304. PMLR, 2018.

[43] A. Beygelzimer, V. Dani, T. Hayes, J. Langford, and B. Zadrozny. Error limiting reductions between classification tasks. In *Proceedings of the 22nd international conference on Machine learning*, pages 49–56, 2005.