Augmented Reality and Proxy Grippers Improve Demonstration-based Robot Skill Learning

Carl Mueller
Univ. of Colorado Boulder
Boulder, Colorado, USA
carl.mueller@colorado.edu

Matthew B. Luebbers
Univ. of Colorado Boulder
Boulder, Colorado, USA
matthew.luebbers@colorado.edu

Aaquib Tabrez
Univ. of Colorado Boulder
Boulder, Colorado, USA
mohd.tabrez@colorado.edu

Bradley Hayes
Univ. of Colorado Boulder
Boulder, Colorado, USA
bradley.hayes@colorado.edu

***Abstract*—Learning from Demonstration (LfD) is a powerful approach that allows novice users to teach robots new skills. However, the quality of user-provided demonstrations is a crucial factor in determining the performance of the learned skill. Poor demonstrations may result in poorly performing learned models, leading to unsafe robot execution and reduced trust. Existing modes of demonstration, such as kinesthetic teaching or teleoperation, have specific advantages and shortcomings that make one preferable over the other depending on the context. For instance, kinesthetic demonstration offers accurate skill representation but can be challenging for untrained users and entirely infeasible for certain robot types or environmental contexts. In contrast, teleoperation does not rely on physical manipulation, expanding the scope of use cases, at the cost of being unintuitive for complex, high degree-of-freedom robots, potentially resulting in poor demonstrations. To address this tradeoff, we propose a novel demonstration input method, extending the recently proposed instrumented tongs technique, wherein a tracked pair of tongs is used by the human demonstrator to serve as a proxy for a robotic end-effector. We combine this method with an augmented reality (AR) interface to visualize and obtain live assessment on what the robot is learning from the provided demonstration, in essence introducing the real-time feedback benefits of physically manipulating a robot to an input method which does not suffer from the ergonomic and feasibility drawbacks of kinesthetic demonstration. We provide a detailed description of our novel demonstration input system and its intended capabilities. Finally, we propose a human-subjects study to evaluate the effectiveness of our method on practically grounded robotic applications, such as mailbox delivery, glue tracing, and stacking.**

***Index Terms*—Learning from Demonstration, Augmented Reality, Robot Learning, Human-Robot Interaction, User Interfaces**

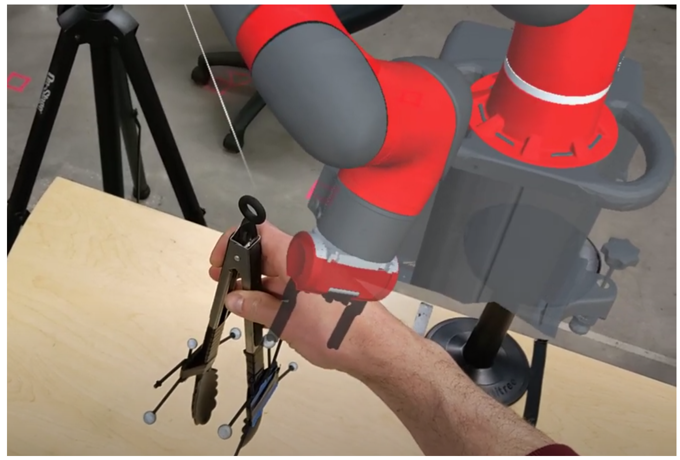

Fig. 1: The ARPOC-LfD system integrates teleoperated instrumented tongs, serving as a surrogate for a robotic end-effector during demonstrations, with an augmented reality-based interface. This interface allows users to visualize in real-time how a virtual robotic agent would execute the demonstration as if it were an actual robot, using online pose optimization to offer self-correcting feedback, thereby enhancing the quality of the demonstration data.

I. Motivation

Learning from Demonstration (LfD) methods enable users to teach robots through demonstration, without requiring programming expertise. This approach provides a flexible and adaptable mechanism for robot control, making it suitable for use in dynamic or human-centric environments [1]. However, the quality of demonstrations provided by users is highly impactful on the performance of learned skills. Existing demonstration methods, such as kinesthetic teaching and teleoperation, have specific benefits and drawbacks that make them more or less appropriate for specific applications [2]. Kinesthetic demonstration excels in skill representation but faces challenges in certain physical situations, while teleoperation bypasses physical interaction but may struggle with complex robots. To address these issues, we propose a novel method that combines the benefits of both kinesthetic and teleoperation-based demonstration methods by leveraging instrumented tongs (a teleoperation device serving as a proxy for the robot end-effector) and an AR-based interface. We call this method Augmented Reality-based Pose Optimization for Constrained Learning from Demonstration (ARPOC-LfD).

The mode of interaction chosen for an LfD system plays a critical role in capturing the intended goals of the human demonstrator, affecting the quality of the demonstration and the fidelity of the learned models. Furthermore, for successful human-robot collaboration, robots must be able to adapt to changes in the environment and incorporate human preferences when necessary [3], [4]. Making these input methods more accessible and user-friendly to end-users is necessary to ensure the safe and effective deployment of robots in real-

This work was funded by the Army Research Lab STRONG Program (#W911NF-20-2-0083) and National Science Foundation (award #1830686).

world environments [5], [6]. By improving the quality of demonstrations and enhancing the robot's ability to adapt to changes in the environment, LfD methods can enable more robust, generalizable, and safer learning models.

Towards these end goals, ARPOC-LfD enables users to generate demonstration trajectories using instrumented tongs as a stand-in for the robotic end-effector without the traditional uncertainty introduced by the correspondence problem, while AR holograms of the robot are closely tracked within the user's visual field (see Figure 1). This approach overcomes challenges in teleoperation by providing real-time, in-environment visualization of how a virtual robotic agent would track the end-effector, similar to kinesthetic demonstration. Through AR visualization, ARPOC-LfD improves the transparency and adaptability of teleoperation devices such as instrumented tongs, allowing users to assess the quality of a provided demonstration, accepting or rejecting demonstration data after replaying it post-task.

The core contributions of this work are as follows:

- An online feasible pose-optimization system that targets user-supplied end-effector poses provided by an instrumented tong device to generate real-time pareto-optimal robot configurations.
- An augmented-reality-based interface that provides users with a visual hologram of a robotic agent that tracks the instrumented tongs as they provide a demonstration in real-time.
- A post-task AR visualization of demonstrated trajectories that allows users to accept or reject the demonstrations.
- A proposed human-subject study evaluation of the ARPOC-LfD system involving real-world, practically grounded robotic manipulation applications.

II. Related Work

A. Robot Learning from Demonstration

Learning from Demonstration (LfD) refers to a collection of methods used to teach robots how to perform tasks by observing human input [7]. The goal of demonstration is to convey the nature of a skill to a robot in such a way that the learned model closely resembles the latent ground truth model held by the demonstrator [1]. In LfD applications, a human user typically interacts with a robotic system by providing a demonstration through one of many possible input modes, such as kinesthetic demonstration. While the mode of demonstration can vary, it is essential that the robot's learned behavior aligns with the intended task. After the user provides the skill demonstration, the system learns the intended behavior or model through one of the three broad classes of learning and representation methods: plan learning, functional optimization, or policy learning [8].

There are various modalities of human-robot interaction (HRI) for robot skill learning, including passive observation, teleoperation, and kinesthetic demonstration. Kinesthetic demonstration is a popular and widely used method of LfD systems [2], [7], [9]. When performed correctly, kinesthetic demonstration can provide a precise and accurate representation of the skill being learned [9], [10]. However, kinesthetic demonstration may be challenging for novice users who lack training, resulting in sub-optimal demonstrations [2]. Additionally, there are many instances where kinesthetic demonstration is not feasible (e.g., if the robot is large or otherwise incapable of being physically manipulated, or if the operational environment is too dangerous for humans)

Teleoperation is an alternative approach to kinesthetic demonstration, which does not require physical manipulation of the robotic agent but instead relies on a controlling device [1], [11]. Teleoperation may not always provide an intuitive means of controlling the robotic agent, especially for complicated high-degree-of-freedom robot agents. A difficult-to-use controller can result in poor agent behavior, as the user cannot easily overcome the poor mapping of teleoperation inputs to control outputs. In the context of LfD, this difficulty can lead to poor demonstrations [2], [8] or be cognitively burdensome to the operator during demonstration [12].

One approach to improving the effectiveness of teleoperation is to use a proxy device that creates a more intuitive mapping between the controller and the agent. For example, Fang et al. [13] used a data glove as a proxy for the end-effector, providing target points for the robot as a set of demonstrations. Similarly, Praveena et al. [2] introduced instrumented tongs that serve as a proxy for the end-effector. These methods are intuitive for users as they use their own hand to demonstrate, rather than physically manipulating the robotic arm as in kinesthetic demonstration or attempting to learn how to use a difficult controller.

The challenge with using data gloves or instrumented tongs as a proxy for the end-effector is that control over the degrees of freedom of the robot is no longer provided by the user, making reliance on kinematics equations crucial for producing agent configurations [2]. This approach only supplies pose targets, which may present challenges towards generating feasible agent configurations, especially given the redundancy of high degree-of-freedom manipulator arms. ARPOC-LfD addresses this issue by utilizing the online pose-optimization framework CollisionIK [14] to generate robot configurations in real time that maintain feasibility and avoid collision states. During a demonstration, users generate pose targets, and the optimization engine produces configurations that satisfy the pose target, avoid collisions and joint limit constraints, and comply with any included constraint terms. The integration of task-space constraint terms allows for configurations that closely adhere to task-specific constraints, improving the quality of demonstrations.

B. Augmented Reality within Robotics

AR interfaces are effective for scenarios where in-situ visualization within the environment is desired [15]. Research has demonstrated that augmented reality interfaces can significantly enhance human-robot collaboration by facilitating new methods of enhancing robotic control [16]–[18], enabling safe movement in shared spaces [19], [20], and promoting effective teamwork through the communication of robot knowledge and intended plans [3], [21], [22]. This motivated us to utilize an AR interface in combination with instrumented tongs to address some of the challenges associated with teleoperated proxy devices. Specifically, ARPOC-LfD utilizes augmented reality (AR) to achieve two primary functionalities: 1) real-time holographic visualization of the robotic agent, which tracks with the instrumented device as users provide demonstrations to enable users to self-correct their demonstration, and 2) a replay option that enables users to replay the demonstrations and accept or reject them as needed.

The inspiration for using AR as a mechanism for self-correction in ARPOC-LfD was drawn from prior research in psychology and HRI. Contrary to the notion that error avoidance should be the ultimate goal of learning, studies by Metcalfe et al. indicate that errorful learning followed by corrective feedback is more beneficial to both students and teachers alike [23]. Moreover, Freedberg et al. describe a study that suggests both positive and negative feedback are helpful for the learning process, but negative feedback may be more advantageous [24]. Within the realm of human-robot interaction, the use of visualization as a feedback mechanism has been shown to enhance collaboration [17]–[19], [25].

The second functionality of ARPOC-LfD, which enables users to replay their demonstrations, is inspired by the work of Luebbers et al. [5]. In their work, they developed an augmented reality (AR) system called ARC-LfD for constrained learning from demonstration, which allows users to maintain, update, and adapt learned skills through in-situ visualizations. The ARC-LfD system enables users to examine a sample trajectory from a learned skill visualized in AR through an overlay in the workspace environment. This skill visualization improves safety by enabling operators to preview robot behavior without the need for actual skill execution [26].

III. System Design

The ARPOC-LfD system consists of two components: a hardware stack and a software stack. The hardware stack involves the use of instrumented tongs, which serve as a proxy for the end-effector during demonstrations (see Figure 2). The software stack is composed of two subsystems (see Figure 3) that communicate through the Robot Operating System (ROS). The first subsystem is an online feasible pose-optimization subsystem that generates feasible robot configurations in real-time. The second subsystem is an AR subsystem that provides live, real-time visualization of a robot hologram as it tracks the instrumented tongs during demonstrations.

A. Instrumented Tongs

Our design for the instrumented tongs was inspired by Praveena et al. [2]. In their study, Praveena et al. introduced a novel input method for users to provide demonstrations, taking inspiration from kitchen tongs. They observed that despite the clumsy form of pinch grasp of kitchen tongs, people could adeptly use them to perform a wide range of manipulations (e.g., serving foods like spaghetti). The instrumented tongs act as a proxy for the robot gripper's end-effector while providing demonstrations. They compared the instrumented tongs with more traditional forms of input methods, such as free-hand manipulation, kinesthetic guidance, and teleoperation, and found that the instrumented tongs provide high-quality demonstrations and a positive experience for the demonstrator while offering good correspondence to the target robot.

In our design, we track the instrumented tongs using an OptiTrack motion capture system via an infrared marker, as illustrated in Figure 2. We use two separate groups of markers to enable the system to detect a closed gripper based on the center-point distance between each group.

Fig. 2: Instrumented tongs utilized in the ARPOC-LfD system. The tongs are tracked using the OptiTrack motion capture system via infrared markers. The two distinct groups of markers enable the system to detect a closed gripper based on the center-point distance between the groups.

B. Online Feasible Pose-optimization Subsystem

The pose optimization subsystem utilizes a multi-objective non-linear constrained optimization program to generate real-time pareto-optimal configurations. This optimization engine is an extension of CollisionIK, a Rust-based software developed by Rakita et el. [14]. Task constraints are incorporated as terms in the multi-objective function, allowing for the integration of multiple constraints alongside collision avoidance, self-collision avoidance, and joint limit constraints implemented in the CollisionIK software.

C. Augmented-reality Subsystem

The second subsystem within the ARPOC-LfD software stack is designed to enable real-time visualization of a robot hologram tracking the instrumented tongs during demonstrations, utilizing an AR headset. Specifically, the Microsoft HoloLens 2 is employed to display a rendered robot agent, such as a Sawyer robotic arm, as a hologram to users. These visualizations are aligned such that the end-effector of the holographic agent coincides with the end-point of the

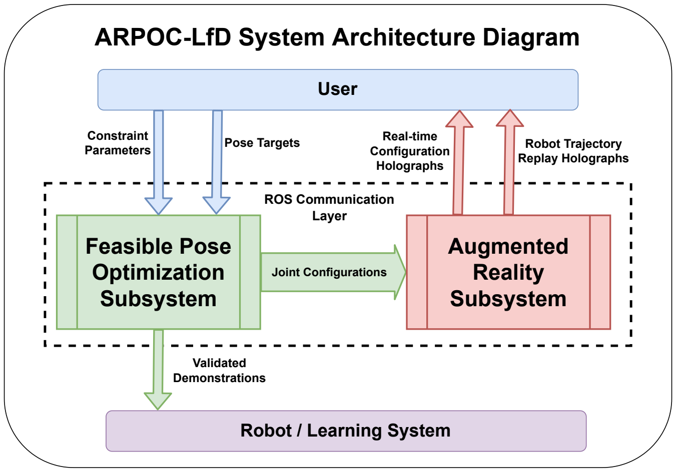

Fig. 3: A schematic of the ARC-LfD system architecture: The user (blue) provides demonstrations using an instrumented tong, along with optional constraint parameters. Through a Robot Operating System communication layer, the feasible pose optimization subsystem (green) delivers real-time joint configurations to the AR subsystem (red). The AR subsystem enables users to view the real-time holographic robot configuration and replay the demonstrated trajectory post-task for acceptance or rejection. Once demonstrations are approved by users, they are sent to the robot learning system (purple).

instrumented tongs (see Figure 1). The tongs are tracked via the OptiTrack Motion Capture system, which produces pose targets for the optimizer based on the position and orientation of the tongs in the environment. The system can also determine whether the tongs are open or closed, based on the distance between the track position of each side of the tong.

To provide visual feedback on tracking error, the system will utilize a color-changing robot, indicating any discrepancies between the position of the tong and the resulting end-effector position. This feedback is based on a threshold distance, determined by the forward kinematics of the optimizer-generated robot configuration. Furthermore, ARPOC-LfD includes a feature that enables users to replay demonstrations and accept or reject them as needed.

ARPOC-LfD employs two feedback mechanisms to facilitate self-correction in human demonstrators. The primary feedback mechanism involves holographic visuals of a robot agent tracking pose targets provided by an instrumented tongs device. The second feedback mechanism provides an indication of the divergence between the virtual holographic end-effector and the physical target of the tong. The greater the tracking error, the more pronounced the indication that the user is providing an unobtainable pose target.

D. Interaction Flow

The interaction flow of ARPOC-LfD aims to enable users to generate demonstration trajectories of robot configurations produced by the online pose optimization subsystem, as illustrated in Figure 4. Initially, users supply pose targets using instrumented tongs tracked by the OptiTrack motion capture system (Step 1, Figure 4). These pose targets are then fed into the feasible pose-optimization engine, which generates a Pareto-optimal robot configuration that best targets the provided pose, while optimizing for static collision avoidance, self-collision avoidance, dynamic obstacle avoidance, and task space constraints (Step 2, Figure 4). If the user is currently in the midst of a demonstration, configurations are transmitted to the AR subsystem, which then generates holograms of the robot agent in the user's visual field via Microsoft Hololens (Step 3, Figure 4). Upon completion of a potentially desirable demonstration, users are able to replay the entire configuration space trajectory on demand (Step 4, Figure 4).

The final step in the interaction with ARPOC-LfD involves using the visual feedback of the latest robot hologram or the trajectory as a whole to ensure that the demonstrated behavior matches the intended action. One of the significant challenges in using devices as proxies for teleoperation is the lack of control over the robot's degrees of freedom. Additionally, instrumented tongs may not provide immediate insight into the robot agent's capabilities, such as reachability, or whether the robot will execute the intended behavior correctly. ARPOC-LfD addresses these limitations through its real-time visualization feature, which enables users to observe a robot hologram tracking the instrumented tongs during demonstrations. This real-time visualization provides users with insights into how their demonstrations would track if an actual physical robot were present. Additionally, the system's ability to replay demonstrations empowers users to accept or reject a demonstration if it is deemed inconsistent with their intended action.

IV. Proposed System Evaluation

This section outlines the proposed human-subjects study that aims to evaluate the effectiveness of the ARPOC-LfD system in three real-world grounded task scenarios for robot manipulators.

A. Experimental Design

The proposed study will use a 3x1 between-subjects design, where users will be randomly assigned to one of three conditions: 1) kinesthetic demonstration, 2) instrumented tongs demonstration without AR visualization, and 3) instrumented tongs demonstration with AR visualization (full-stack ARPOC-LfD system). A between-subjects design was selected to obtain more objective results on the effectiveness of each form of demonstration and to avoid learning effects. Participants in each condition will be asked to demonstrate three tasks: a mailbox delivery task, a glue tracing task, and a stacking task, with at least three demonstrations for each task.

B. Experimental Conditions

Condition 1 - Kinesthetic Demonstration: Participants will physically move a robotic manufacturing arm (a Rethink Robotics Sawyer) through the intended skill, tracing out the trajectory the participant is attempting to teach the robot.

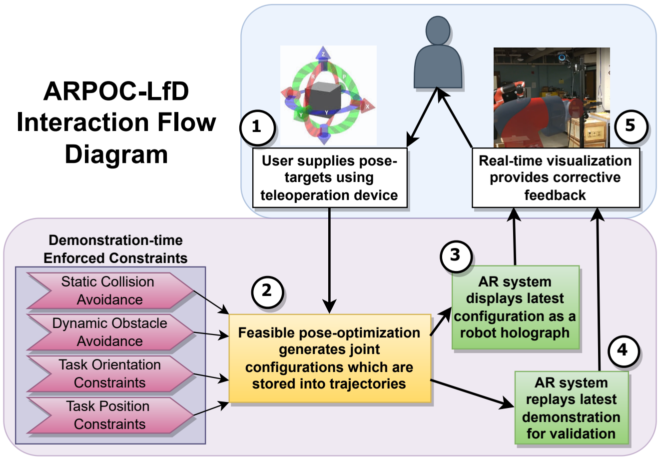

Fig. 4: The ARPOC-LfD interaction flow diagram. The bottom half indicates system processes whereas the top half indicates how the user interacts with the system. ARPOC-LfD enables users to input pose targets using instrumented tongs (1) that are processed by a feasible pose-optimization engine (2) to generate a Pareto-optimal robot configuration. This configuration is transmitted to the AR subsystem to create robot agent holograms (3). Users can confirm that demonstrated behavior aligns with intended action using visual feedback from the latest robot hologram (5) or trajectory replay (4), made possible by the system's real-time visualization feature.

Condition 2 - Instrumented Tongs Demonstration, no AR: Participants will use the instrumented tongs device to mimic the robot's end-effector and manipulate objects for the intended skill, similar to Praveena et al. [2]. The instrumented tongs resemble an ordinary pair of kitchen tongs, with minimal sensors/motion tracking apparatus attached to it. The position and rotation of the device will be continuously tracked throughout the demonstration to provide trajectory data using OptiTrack. The physical Sawyer robot will not be present in this condition - participants will only interact with the instrumented tongs device.

Condition 3 - Instrumented Tongs Demonstration with AR Visualization: This condition involves the same input device as Condition 2, but with the addition of an augmented reality interface. Participants will wear a HoloLens 2 headset while demonstrating tasks using the instrumented tongs. The headset will display a virtual 3D model of Sawyer following the instrumented tongs with its arm in real-time, providing visualization of how Sawyer would perform the task. After providing a demonstration, participants will be able to replay the demonstration by issuing a command with their headsets to accept or reject the provided demonstration.

C. Experimental Tasks

Participants will be asked to provide a minimum of three demonstrations each for three tasks. These tasks are designed to be challenging in specific way, which may make one form of demonstration easier than the others. Some tasks may favor the precision of kinesthetic demonstration, while others may favor the convenience and ease of using the instrumented tongs. The purpose of these tasks is to provide a diverse set of scenarios that can effectively evaluate the different forms of demonstration.

Task I - Mailbox Delivery: The goal of this task is for participants to demonstrate delivery of an object to a mailbox (see Figure 5a). The task consists of four subtasks: 1) opening the mailbox door, 2) picking up a small block, 3) placing the block inside the mailbox, and 4) closing the mailbox door. This task represents a real-world scenario where objects need to be manipulated and placed in a confined space.

The opening of the mailbox is a particularly challenging movement to demonstrate, especially using kinesthetic demonstration. Precise control of the end-effector is required to grasp the handle and to open and close the door.

Task II - Glue Tracing: Here, participants are required to demonstrate the tracing of a dummy glue stick around the perimeter of an object, (see Figure 5b). The following subtasks are required for a successful demonstration: 1) picking up the glue stick from a jar, 2) avoiding an unregistered collision object (it will not be included in any collision avoidance mechanism), 3) tracing around the object on the workbench and, 4) returning the glue stick back to its receptacle.

This task is likely to be more advantageous for kinesthetic demonstration as the instrumented tongs approach might not have the fidelity needed to successfully provide a demonstration.

Task III - Stacking: In this task, participants will be demonstrating the stacking of multiple objects in the center of the workspace environment (see Figure 5c). The task requires the participants to pick up each of four colored blocks in a specified order and stack them on top of each other.

The blocks are placed at the corners of the workspace near the robot's operating space limits, which emphasizes the importance of demonstrating in a manner that the robot can feasibly reach the blocks.

D. Measurement and Evaluation

Here we describe the various objective and subjective measures used to assess user experience and the quality of provided demonstrations to evaluate ARPOC-LfD against the two baseline conditions. There are three objective measures: demonstration trajectory mean warping distance, physical robot feasibility percentage, and task execution percentage.

Demonstration Trajectory Mean Warping Distance: To evaluate each user's set of demonstrations, we will calculate the mean warping distance. We will generate a representative trajectory of the set using the Gaussian Mean Regression (GMR) methodology [9], [27], [28]. This representative trajectory will serve as the candidate trajectory for the Dynamic Time Warping distance measure. We will use this distance measure to calculate the mean and variance of the set of trajectories warped against the GMR-produced candidate.

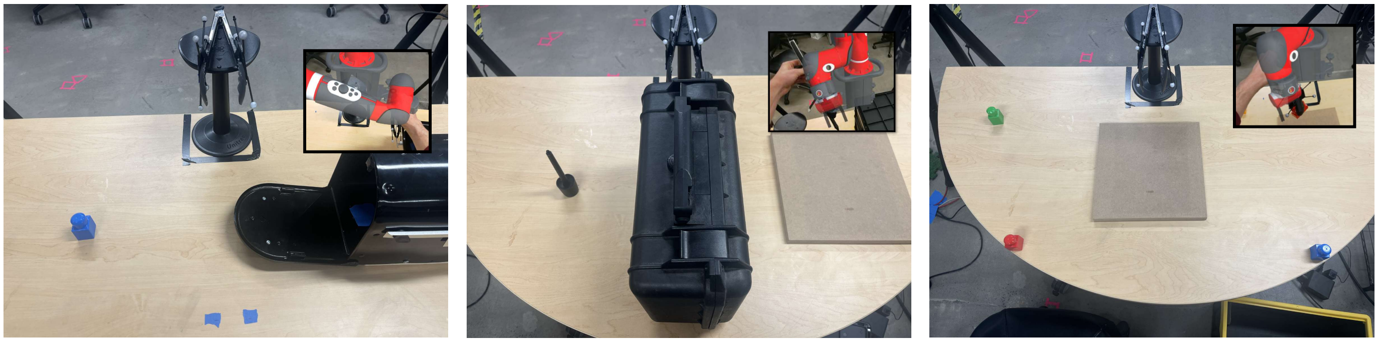

(a) (b) (c)

Fig. 5: Our proposed evaluation tasks include: a) Mailbox Delivery - participants showcase opening a mailbox, inserting an object, and closing it, b) Glue Tracing - participants show tracing with a dummy glue stick around an object's perimeter, and c) Stacking - participants exhibit stacking multiple objects at the center of the workspace environment.

Physical Robot Feasibility Percentage: We will test each user's set of demonstrations for robot execution feasibility and calculate the success rate as a percentage.

Task Execution Percentage: We will test the ability of demonstration trajectories to complete the task successfully when executed on a physical agent and report the success rate as a percentage.

We will also use two broad classes of subjective measures to evaluate and compare the demonstration trajectories from each condition: expert demonstration analysis and subjective questionnaires.

Expert Demonstration Analysis: Each user's set of demonstrations will undergo analysis by multiple non-study-affiliated robotics experts who will rank them on a scale of 1 to 10 based on their belief that the trajectory is a proper demonstration for the task.

User Questionnaires: During the experiment, each study participant will complete both between-task and post-experiment surveys. The post-experiment surveys will incorporate 7-point Likert-scale items derived from established questionnaires in the robotics and explainable AI community. These surveys are geared towards assessing reliability and usability (SUS) [29], trust and confidence [30], and explainability [31]. Meanwhile, the between-task surveys will utilize the NASA Task Load Index assessment [32], along with a trust/confidence assessment.

V. Conclusion

In this work, we introduce a new input method for Learning from Demonstration applications called Augmented Reality-based Pose Optimization for Constrained Learning from Demonstration (ARPOC-LfD). This method aims to combine the ergonomics and broad usability of teleoperation-based demonstration with the accurate skill representation and live feedback from a physically embodied robot found in kinesthetic-based demonstration. ARPOC-LfD uses instrumented tongs as a proxy for a robotic end-effector during demonstration, while the real-time feasibility of joint configurations and the demonstrated trajectory is visualized through an augmented reality interface. We provide a detailed description of each component of this proposed system. Finally, we propose a human-subjects user study to evaluate the effectiveness of our novel demonstration input interface on practically grounded real-world robotic manipulation applications, such as mailbox delivery, glue tracing, and stacking tasks.

References

[1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.

[2] P. Praveena, G. Subramani, B. Mutlu, and M. Gleicher, "Characterizing input methods for human-to-robot demonstrations," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 344–353.

[3] A. Tabrez, M. B. Luebbers, and B. Hayes, "A survey of mental modeling techniques in human–robot teaming," *Current Robotics Reports*, vol. 1, no. 4, pp. 259–267, 2020.

[4] B. Hayes and B. Scassellati, "Challenges in shared-environment human-robot collaboration," *learning*, vol. 8, no. 9, 2013.

[5] M. B. Luebbers, C. Brooks, C. L. 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 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3794–3800.

[6] A. Tabrez, J. Kawell, and B. Hayes, "Asking the right questions: Facilitating semantic constraint specification for robot skill learning and repair," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6217–6224.

[7] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *ICML*, vol. 97. Citeseer, 1997, pp. 12–20.

[8] 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*, vol. 3, 2020.

[9] B. Akgun, M. Cakmak, K. Jiang, and A. L. Thomaz, "Keyframe-based learning from demonstration," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 343–355, 2012.

[10] B. Akguen, K. Subramanian, and A. L. Thomaz, "Novel interaction strategies for learning from teleoperation." in *AAAI Fall Symposium: Robots Learning Interactively from Human Teachers*, vol. 12, 2012, p. 07.

[11] P. Bakker, Y. Kuniyoshi *et al.*, "Robot see, robot do: An overview of robot imitation," in *AISB96 Workshop on Learning in Robots and Animals*, vol. 5. Citeseer, 1996.

[12] I. Havoutis and S. Calinon, "Learning from demonstration for semi-autonomous teleoperation," *Autonomous Robots*, vol. 43, no. 3, pp. 713–726, 2019.

[13] B. Fang, D. Guo, F. Sun, H. Liu, and Y. Wu, "A robotic hand-arm teleoperation system using human arm/hand with a novel data glove," in *2015 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2015, pp. 2483–2488.

[14] D. Rakita, H. Shi, B. Mutlu, and M. Gleicher, "Collisionik: A per-instant pose optimization method for generating robot motions with environment collision avoidance," *arXiv preprint arXiv:2102.13187*, 2021.

[15] H. Hedayati, M. Walker, and D. Szafir, "Improving collocated robot teleoperation with augmented reality," in *2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2018, pp. 78–86.

[16] J. Weisz, P. K. Allen, A. G. Barszap, and S. S. Joshi, "Assistive grasping with an augmented reality user interface," *The International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 543–562, 2017.

[17] M. E. Walker, H. Hedayati, and D. Szafir, "Robot teleoperation with augmented reality virtual surrogates," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 202–210.

[18] C. Brooks and D. Szafir, "Visualization of intended assistance for acceptance of shared control," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.

[19] M. Walker, H. Hedayati, J. Lee, and D. Szafir, "Communicating robot motion intent with augmented reality," in *2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2018, pp. 316–324.

[20] E. Rosen, D. Whitney, E. Phillips, G. Chien, J. Tompkin, G. Konidaris, and S. Tellex, "Communicating robot arm motion intent through mixed reality head-mounted displays," in *Robotics Research*. Springer, 2020, pp. 301–316.

[21] M. Diehl, A. Plopski, H. Kato, and K. Ramirez-Amaro, "Augmented reality interface to verify robot learning," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, pp. 378–383.

[22] 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*, 2022, pp. 1256–1264.

[23] J. Metcalfe, "Learning from errors." *Grantee Submission*, vol. 68, pp. 465–489, 2017.

[24] M. Freedberg, B. Glass, J. V. Filoteo, E. Hazeltine, and W. T. Maddox, "Comparing the effects of positive and negative feedback in information-integration category learning," *Memory & cognition*, vol. 45, no. 1, pp. 12–25, 2017.

[25] D. Szafir, "Mediating human-robot interactions with virtual, augmented, and mixed reality," in *2019 International Conference on Human-Computer Interaction (HCI)*. Springer, 2019, pp. 124–149.

[26] D. Krupke, F. Steinicke, P. Lubos, Y. Jonetzko, M. Görner, and J. Zhang, "Comparison of multimodal heading and pointing gestures for co-located mixed reality human-robot interaction," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–9.

[27] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.

[28] S. Calinon, F. Guenter, and A. Billard, "On learning, representing, and generalizing a task in a humanoid robot," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 37, no. 2, pp. 286–298, 2007.

[29] J. Brooke *et al.*, "Sus-a quick and dirty usability scale," *Usability evaluation in industry*, vol. 189, no. 194, pp. 4–7, 1996.

[30] M. Lewis, K. Sycara, and P. Walker, "The role of trust in human-robot interaction," in *Foundations of trusted autonomy*. Springer, Cham, 2018, pp. 135–159.

[31] R. R. Hoffman, S. T. Mueller, G. Klein, and J. Litman, "Metrics for explainable ai: Challenges and prospects," *arXiv preprint arXiv:1812.04608*, 2018.

[32] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.