

Enhancing Incremental-learning process of Robotic with Augmented Reality

Chao Wang, Anna Belardinelli, Daniel Tanneberg, Stephan Hasler, Theodoros Stouraitis, Michael Gienger

2023

Preprint:

Copyright ACM 2023. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in CHI 2023,
<https://doi.org/https://doi.org/10.1145/3544549.3583889>.

Explainable Human-Robot Training and Cooperation with Augmented Reality

Chao Wang*
 chao.wang@honda-ri.de
 Honda Research Institute Europe GmbH
 Offenbach, Hessen, Germany

Anna Belardinelli*
 anna.belardinelli@honda-ri.de
 Honda Research Institute Europe GmbH
 Offenbach, Hessen, Germany

Stephan Hasler*
 stephan.hasler@honda-ri.de
 Honda Research Institute Europe GmbH
 Offenbach, Hessen, Germany

Theodoros Stouraitis*
 theodoros.stouraitis@honda-ri.de
 Honda Research Institute Europe GmbH
 Offenbach, Hessen, Germany

Daniel Tanneberg*
 daniel.tanneberg@honda-ri.de
 Honda Research Institute Europe GmbH
 Offenbach, Hessen, Germany

Michael Gienger*
 michael.gienger@honda-ri.de
 Honda Research Institute Europe GmbH
 Offenbach, Hessen, Germany

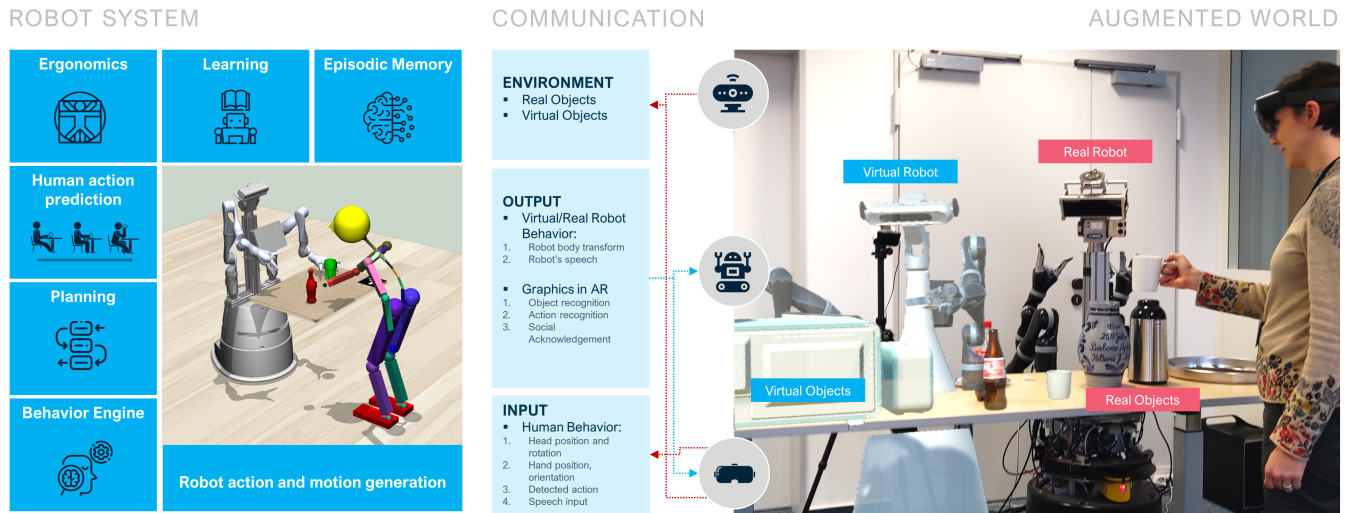


Figure 1: System architecture for AR-mediated training and cooperation with robots.

ABSTRACT

The current spread of social and assistive robotics applications is increasingly highlighting the need for robots that can be easily taught and interacted with, even by users with no technical background. Still, it is often difficult to grasp what such robots know or to assess if a correct representation of the task is being formed. Augmented Reality (AR) has the potential to bridge this gap. We demonstrate three use cases where AR design elements enhance the explainability and efficiency of human-robot interaction: 1) a

human teaching a robot some simple kitchen tasks by demonstration, 2) the robot showing its plan for solving novel tasks in AR to a human for validation, and 3) a robot communicating its intentions via AR while assisting people with limited mobility during daily activities.

CCS CONCEPTS

• **Human-centered computing** → **User interface toolkits.**

KEYWORDS

explainability, human-robot interaction, augmented reality

ACM Reference Format:

Chao Wang, Anna Belardinelli, Stephan Hasler, Theodoros Stouraitis, Daniel Tanneberg, and Michael Gienger. 2023. Explainable Human-Robot Training and Cooperation with Augmented Reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3544549.3583889>

*All authors contributed equally to this research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
 CHI EA '23, April 23–28, 2023, Hamburg, Germany
 © 2023 Copyright held by the owner/author(s).
 ACM ISBN 978-1-4503-9422-2/23/04.
<https://doi.org/10.1145/3544549.3583889>

1 INTRODUCTION

Among the foremost conditions to have socially assistive robots enter our homes and act around us, is the possibility even for non-expert users to be able to intuitively interact with the robot and to teach it new tasks. Such systems are meant for long-term interaction, hence they will need to be flexible enough to handle a large number of tasks and to learn new ones. This means that their functioning will be rather complex, relying on multiple AI and machine learning algorithms which cannot be fathomed by the end user. At the same time, robots will need to inform the user about their workings in an effective way, so that the user can adjust its mental model and understand what to expect from the robot without the need for specific training, lengthy instructions, or repeated observations of the robot. Such a feature would go a long way in improving acceptance and trust in robotic assistants by the public. This need, already emerged in recent years in disembodied AI systems and tackled in the eXplainable AI (XAI) field [1, 4, 15, 16], has been previously denoted as interpretability or legibility in Mobile Robotics [3], and has been recently characterized by specific robotic dimensions [5]. Indeed, while pure AI systems usually provide an a-posteriori explanation on decisions/predictions made by some machine learning algorithm, robots are embodied agents, autonomously acting in the real world and possibly interacting and cooperating with humans. This situatedness and agency calls for different kinds of explanations, not just regarding past decisions, but more critically about current assessments, future plans, and intentions. Such information needs to be timely communicated to the human user and interactively negotiated with them in a quick and unambiguous fashion. In Human-Robot Interaction (HRI) similar interactions have often been managed by imitating interactions and social cues exchanged between humans, e.g., gaze contact and shared attention, speech, gestures [14], with the idea that behavioral human-likeness can improve intuitiveness. Still, due to their familiarity, these cues might falsely suggest that the robot has similar perceptual and understanding capabilities as the human counterpart. In this sense, other interaction modalities might integrate these cues and offer an insight into the robot's mind that helps bridging the user's mental model gaps.

We propose a concept utilizing AR to improve the user experience of interacting with robot assistants by displaying intuitive explanatory hints about the robot perception, learning, generalization, and planning at different levels. While Augmented/Mixed Reality solutions have been recently increasingly spreading in robotics [2, 7, 20], these have typically targeted debugging by trained experts [9], visualizing internal states and movement intentions [8, 12], teleoperation [21], and simplifying learning by demonstration of movement trajectories [6]. Here, we introduce our robotic system which leverages AR capabilities across different typical use cases in robotics: learning new tasks by demonstration, devising a plan for a new problem, and assisting a human during a physical task in an ergonomic way. By exploiting real and virtual objects and by integrating human-like social cues with explanatory virtual design elements [13], this framework showcases how a hybrid workspace can be shared between humans and robots. Such interfaces enhance the interaction by making the robot's state of mind

visually apparent and transparent to the user, grounding it on the current human perception of the scene.

2 SYSTEM ARCHITECTURE

Our system is composed of modules realizing the back-end functionalities (learning, planning, prediction and motion generation), and the front-end interface capabilities (visualization and interaction) of the AR glasses (see Fig. 1).

Back-end robotic system: The physical robot is a torso with two Kinova arms¹ with 7 DOF each and a pan-tilt unit, part of the mobile platform "Johnny" [18, 19]. The robot's cognitive skills are realized in multiple ROS² nodes, communicating with each other and with the front-end interface. The *behavior engine* orchestrates the social behavior of the robot. It controls the robot gaze, while concurrently issuing XAI cues to be shown in the AR environment (cfr. [17]). It also regulates the speech interaction, e.g. acknowledging the user commands or asking curiosity-driven questions (see Sec. 3). The *episodic memory* collects world state observations during demonstrations by the human teachers. The *learning* module realizes symbolic skill learning which integrates demonstrations from the episodic memory into a knowledge graph, and generates new hypotheses to be queried to the user (Sec. 3). The *planning* module uses the learned semantic skills to generate high-level plans to solve new tasks. Such plans are yet to be checked with the human tutor (see Sec. 4). The *human action prediction* module operates when assisting a user. The robot can predict what the user is intending to achieve, and plan a supportive action, e.g., moving the next required object closer to the user. The *ergonomics* module generates interventions in an ergonomically optimal way for the user (see Sec. 5). Finally, the *action generation* module (Fig. 1) receives action commands from other modules (e.g., where to look, what to grasp, etc) and executes the corresponding motor behavior.

Front-end interface: Virtual objects, a virtual robot, and XAI cues (AR graphical elements) are displayed in the mixed-reality environment via the HoloLens³. The HoloLens can scan the surroundings, build up 3D meshes of the environment objects and locate itself in the room, which enables it to stably overlay graphics in the environment considering occlusions with real objects.

The shared environment, accessible to the user through the HoloLens (see Fig. 1, right), includes multiple real and virtual objects. The object poses are continuously sent to the back-end, by the HoloLens for virtual objects and by a static camera for real objects (endowed with fiducial markers). Users can interact with the virtual objects in a similar way as with real ones; they can, for instance, pick up the bread, put it into a slot of the toaster, and push the button.

The HoloLens is also detecting the user's behavior and communicating it to the back-end system. This includes the user's head position, orientation and speech input. More importantly, as the HoloLens can track the user's hand and fingers, then the manual actions (e.g., "pick" or "drop") are also detected. The manipulation of objects by a human hand is implemented via Microsoft MRTK

¹<https://www.kinovarobotics.com/product/gen2-robots>

²<https://www.ros.org>

³<https://www.microsoft.com/en-us/hololens>




Figure 2: Learning from demonstration with XAI: while the user demonstrates a new skill the robot signals in AR which actions and objects it recognizes. Afterwards, it asks a related question while highlighting involved objects.

SDK⁴, which enables the corresponding virtual object to stick to the human’s hand while the "picking/holding" gesture is applied, and release from the hand after a "drop" gesture is detected. Moreover, colliders are attached to the user’s fingers, enabling the teacher, for instance, to press the toaster lever, turn on the power button of the microwave and close the microwave door. Finally, the user behavior and related manipulated object information are sent to the back-end system via ROS (see Fig. 1, "INPUT"). The system can also display and animate a holographic virtual robot, which looks almost identical as the physical robot. The HoloLens receives the robot behavior data (see Fig. 1 "OUTPUT"), including the pose of the virtual robot and the speech commands, and visualizes / speaks them. Furthermore, the back-end triggers the display of the XAI cues which are shown in the AR environment. The next sections introduce three use cases for enhancing human-robot interaction via AR based on this system architecture.

3 USE CASE 1: EXPLAINABLE HUMAN-ROBOT INTERACTION FOR IMITATION LEARNING

Learning by demonstration or imitation learning in robotics aims at enabling humans to teach new skills by physically showing the task as they would do to another person. Major limitations in such an approach have been the capability of the robot to interpret and generalize a specific demonstration, and the difficulty that human users have to understand what constitutes a good demonstration for the robot. We therefore developed a two-stage skill learning concept. In the first stage, the user demonstrates a skill to the robot, which acquires it using semantic skill learning concepts. The learned representation of the skill is formed by symbols that encode preconditions, actions, and effects. In the second stage, the robot takes initiative and asks curious questions about the demonstrated task to the user. Both stages are designed to enhance the user’s mental model of the system using AR and social cues.

Specifically, the user (wearing AR glasses) demonstrates a skill to the robot, e.g., by interacting with the (virtual) objects and giving language explanations (skill labels). The robot follows the teacher’s gaze or looks back at him to show its attention. The teacher is further continuously informed about the robot’s perception: Object and action labels (XAI cues) are popping up whenever the teacher gazes at some object or a manual action is recognized (cfr. Fig. 2 and [17]). The state of the environment, including the objects and

agents, is recorded before and after the demonstrated skill by the episodic memory. These observations are parsed into predefined symbolic representations of the environment, consisting of logical predicates. Together with the skill label, these observations are used to learn the symbolic skill, capturing in which situations it can be applied and what is changed in the environment by executing it.

If skills with similar effect but different objects are learned, the robot can generalize these skills using a predefined object hierarchy. This type-generalization of skills enables to consider novel objects in a new task. This symbolic nature of the skill representation enables the robot to explain the learned knowledge or a future plan in an intuitive way.

While such an imitation learning scenario is intuitive and engaging for the user, it still takes time and effort to demonstrate the skills. To speed up the interactive learning process, the robot has been equipped with the ability to ask curious questions about skills and objects, i.e., creating hypotheses that can be presented to the user for confirmation or rejection. This allows teaching additional knowledge to the system in an interactive and, in comparison to full skill demonstrations, faster way. For example, after seeing a demonstration to learn to use the microwave, where the teacher heated milk, the robot can ask whether a similar object (e.g., water) can be heated in the microwave. The XAI virtual elements in this case would highlight the candidate object and the microwave and, after the user’s answer, acknowledge with red/green particles the negative/positive answer (see Fig. 2, right panels).

4 USE CASE 2: VISUALIZATION OF THE ROBOT PLANNING

Learned skills contain both high-level knowledge about the physical effects of different devices and low-level knowledge on how to operate these devices and manipulate objects. We use a standard symbolic STRIPS planner to combine the skills to solve novel and more complex tasks considering present objects, both real and virtual ones. For example, when asked to 'Prepare an ice tea' the planner might consider the kettle or the microwave to heat some water before putting a tea bag inside and the fridge or some ice cubes to cool it later. Such a high-level solution is extended with the required low-level manipulations, like placing objects, opening doors and pressing buttons.

By learning the skills and applying them to new tasks and environments, the system generalizes from previously observed episodes. As a consequence, the generated plans may not be feasible

⁴<https://learn.microsoft.com/en-us/windows/mixed-reality/mrkt-unity/mrkt2/?view=mrktunity-2021-05>




Figure 3: Visualization of planning phases, from left to right: the user instructs the robot by speech; the robot shows its avatar and digital twins of real objects; the avatar executes the plan; the human gives feedback on the plan. Red boxes show the user’s commands as detected by speech recognition.




Figure 4: Instance (a) occurs in a scenario without assistance, where the human bends to pour into the glass, resulting in an inconvenient posture. In contrast, instances (b) and (c) are two consecutive frames from a scenario with assistance. In (b), the robots predicts that the human will pour into the glass and reveals the outcome of its intervention utilising XAI cues while relocating the glass. In (c), the human pours into the glass in a comfortable posture.

or desired and need validation by a human. Here, we propose an interactive system which can show the plan of the robot to the human via AR glasses before it is executed (see Fig. 3). The user can give a command to the robot via speech. Then the robot will generate a plan to solve the query according to its knowledge at different levels. Before execution, the robot will ask the user to validate the plan in AR. The virtual “avatar” of the robot appears overlaid on the physical body of the robot and real object “shadows” (holographic twins) are displayed in the AR glasses. Then the virtual robot will execute the plan with the virtual objects. In this way, the human can understand the robot reasoning and provide feedback to its plan, approving or correcting it.

5 USE CASE 3: COMMUNICATING ROBOT’S INTENTIONS WHILE ASSISTING USERS

Efficient cooperation between a human and a robot includes a number of challenges. These, among others, are: (i) human behaviour anticipation, (ii) predicting the influence of the robot’s actions onto the human’s actions, and (iii) generating legible robot actions [3] that can be easily understood by the human partner. Here, we study these challenges in the context of manipulation where a robot assists a human to perform a sequence of actions. As a typical example, consider a human (sitting on a wheelchair) serving a beverage, as shown in Figure 4. This scenario requires the human to bend/stretch to reach the glass, hence, a robot assistant could be of aid.

To support the human, the robot ideally predicts the sequence of likely human actions and postures [10], assesses the future human postures [11], computes the desired intervention that improves the physical state of the human, and communicates the effects of the intended interventions to the human ahead of time. To achieve these aims, we propose a model-based optimisation approach that performs a physics-based prediction of the human’s actions and enables the robot to decide which intervention will improve the human posture, hence their ergonomic state, and illustrate them to the human via the XAI cues.

In this use case we consider a table-top scenario, where the location of several objects and the intended high-level task is provided to the system. Such a task can be serving a beverage or a bowl of cereal. Our system is able to: (i) predict the sequence of humans actions, e.g. picking up a bottle and then pouring into a glass. (ii) Evaluating the human posture, e.g. human upper-body configuration while pouring, and decide how to adjust continuous quantities, e.g. change the pose of the glass to improve the human’s upper-body configuration while pouring. (iii) Inform the human using AR (e.g. using holograms to show where objects will be relocated), while performing the assistive action. Utilizing AR to reveal the outcome of a future invention allows the user to understand the robot actions, enabling him to comfortably and fluently perform the task.

6 CONCLUSIONS

Our system aims at providing critical information about the robot, alleviating the cognitive and physical load of the user who can act and interact in a natural way, without constantly monitoring and scrutinizing the robot but correcting it only if necessary. While we target here the specific case of robots assisting in a domestic environment, we implement AR-based design concepts that could cater to understanding and justification needs increasingly present in the interaction with physical intelligent systems and hopefully can inspire the broader CHI community with related solutions for smart homes or other human-centered AI systems. Still, the focus of our work is on facilitating robot interaction for non-experts and on increasing the independence of elderly and people with limited mobility, thus creating an impact on accessibility and inclusiveness in human-machine interaction.

REFERENCES

- [1] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bénéto, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58 (2020), 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [2] Zahraa Bassyouni and Imad H. Elhaji. 2021. Augmented Reality Meets Artificial Intelligence in Robotics: A Systematic Review. *Frontiers in Robotics and AI* 8 (sep 2021). <https://doi.org/10.3389/frobt.2021.724798>
- [3] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 301–308. <https://doi.org/10.1109/HRI.2013.6483603>
- [4] David Gunning, Mark Stefik, Jaesik Choi, Timothy Miller, Simone Stumpf, and Guang-Zhong Yang. 2019. XAI—Explainable artificial intelligence. *Science robotics* 4, 37 (2019), eaay7120.
- [5] Thomas Hellström and Suna Bensch. 2018. Understandable robots - What, Why, and How. *Paladyn, Journal of Behavioral Robotics* 9, 1 (jul 2018), 110–123. <https://doi.org/10.1515/pjbr-2018-0009>
- [6] Matthew B Luebbers, Connor Brooks, Minjae John Kim, Daniel Szafrir, and Bradley Hayes. 2019. Augmented reality interface for constrained learning from demonstration. In *Proceedings of the 2nd International Workshop on Virtual, Augmented and Mixed Reality for HRI (VAM-HRI)*.
- [7] Zhanat Makhataeva and Huseyin Varol. 2020. Augmented Reality for Robotics: A Review. *Robotics* 9, 2 (apr 2020), 21. <https://doi.org/10.3390/robotics9020021>
- [8] Leah Perlmutter, Eric Kernfeld, and Maya Cakmak. 2016. Situated Language Understanding with Human-like and Visualization-Based Transparency. In *Robotics: Science and Systems XII*. Robotics: Science and Systems Foundation. <https://doi.org/10.15607/rss.2016.xii.040>
- [9] Alexandros Rotsidis, Andreas Theodorou, Joanna J Bryson, and Robert H Wortham. 2019. Improving robot transparency: An investigation with mobile augmented reality. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 1–8.
- [10] Theodoros Stouraitis, Iordanis Chatz Nikolaidis, Michael Gienger, and Sethu Vijayakumar. 2020. Online Hybrid Motion Planning for Dyadic Collaborative Manipulation via Bilevel Optimization. *IEEE Transactions on Robotics* 36, 5 (2020), 1452–1471. <https://doi.org/10.1109/TRO.2020.2992987>
- [11] Linda van der Spaa, Michael Gienger, Tamas Bates, and Jens Kober. 2020. Predicting and optimizing ergonomics in physical human-robot cooperation tasks. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1799–1805.
- [12] Michael Walker, Hooman Hedayati, Jennifer Lee, and Daniel Szafrir. 2018. Communicating Robot Motion Intent with Augmented Reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, New York, NY, USA, 316–324. <https://dl.acm.org/doi/10.1145/3171221.3171253>
- [13] Michael Walker, Thao Phung, Tathagata Chakraborti, Tom Williams, and Daniel Szafrir. 2022. Virtual, augmented, and mixed reality for human-robot interaction: A survey and virtual design element taxonomy. *arXiv preprint arXiv:2202.11249* (2022).
- [14] Sebastian Walkötter, Silvia Tulli, Ginevra Castellano, Ana Paiva, and Mohamed Chetouani. 2021. Explainable Embodied Agents Through Social Cues. *ACM Transactions on Human-Robot Interaction* 10, 3 (sep 2021), 1–24. <https://doi.org/10.1145/3457188>
- [15] Chao Wang and Pengcheng An. 2021. Explainability via Interactivity? Supporting Nonexperts' Sensemaking of Pre-Trained CNN by Interacting with Their Daily Surroundings. In *Extended Abstracts of the 2021 Annual Symposium on Computer-Human Interaction in Play (Virtual Event, Austria) (CHI PLAY '21)*. Association for Computing Machinery, New York, NY, USA, 274–279. <https://doi.org/10.1145/3450337.3483487>
- [16] Chao Wang and Pengcheng An. 2021. A Mobile Tool That Helps Nonexperts Make Sense of Pretrained CNN by Interacting with Their Daily Surroundings. In *Adjunct Publication of the 23rd International Conference on Mobile Human-Computer Interaction (Toulouse & Virtual, France) (MobileHCI '21)*. Association for Computing Machinery, New York, NY, USA, Article 19, 5 pages. <https://doi.org/10.1145/3447527.3474873>
- [17] Chao Wang and Anna Belardinelli. 2022. Investigating explainable human-robot interaction with augmented reality. In *Proceedings of the 5th International Workshop on Virtual, Augmented and Mixed Reality for HRI (VAM-HRI)*.
- [18] Chao Wang, Stephan Hasler, Manuel Mühlhig, Frank Joublin, Antonello Ceravola, Joerg Deigmoeller, and Lydia Fischer. 2020. Designing Interaction for Multi-Agent System in an Office Environment. In *The Eighth International Workshop of Chinese CHI (Honolulu, HI, USA) (Chinese CHI 2020)*. Association for Computing Machinery, New York, NY, USA, 86–88. <https://doi.org/10.1145/3403676.3404349>
- [19] Chao Wang, Stephan Hasler, Manuel Mühlhig, Frank Joublin, Antonello Ceravola, Joerg Deigmoeller, Lydia Fischer, and Pengcheng An. 2021. Designing Interaction for Multi-Agent Cooperative System in an Office Environment. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (Boulder, CO, USA) (HRI '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 668–669. <https://doi.org/10.1145/3434074.3447277>
- [20] Tom Williams, Daniel Szafrir, Tathagata Chakraborti, and Heni Ben Amor. 2018. Virtual, augmented, and mixed reality for human-robot interaction. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 403–404.
- [21] A.W.W. Yew, S.K. Ong, and A.Y.C. Nee. 2017. Immersive Augmented Reality Environment for the Teleoperation of Maintenance Robots. *Procedia CIRP* 61 (2017), 305–310. <https://doi.org/10.1016/j.procir.2016.11.183>