The Reality-Virtuality Interaction Cube

Tom Williams
Colorado School of Mines
Golden, CO, USA
twilliams@mines.edu

Daniel Szafir
University of Colorado Boulder
Boulder, CO, USA
daniel.szafir@colorado.edu

Tathagata Chakraborti
IBM Research AI
Cambridge, MA, USA
tathagata.chakraborti1@ibm.com

***Abstract*—There has recently been an explosion of work in the humanrobot interaction (HRI) community on the use of mixed, augmented, and virtual reality. In this paper, we present a novel conceptual framework to characterize and cluster work in this new area and identify gaps for future research. We begin by introducing the Plane of Interaction: a framework for characterizing interactive technologies in a 2D space informed by the Model-View-Controller design pattern. We then describe how Interaction Design Elements that contribute to the interactivity of a technology can be characterized within this space and present a taxonomy of Mixed-Reality Interaction Design Elements. We then discuss how these elements may be rendered onto both reality- and virtuality-based environments using a variety of hardware devices and introduce the *Reality-Virtuality Interaction Cube*: a three-dimensional continuum representing the design space of interactive technologies formed by combining the Plane of Interaction with the Reality-Virtuality Continuum. Finally, we demonstrate the feasibility and utility of this framework by clustering and analyzing the set of papers presented at the recent 2018 VAM-HRI Workshop.**

***Index Terms*—Mixed Reality, Augmented Reality, Virtual Reality, Robotics, Human-Robot Interaction.**

I. Introduction

Although research on augmented reality [7], [8], [12], [46], [56] and virtual reality [13], [40] have been steadily progressing over the past several decades, there has been relatively little work using augmented reality (AR) technologies to facilitate human-robot interactions (despite a number of papers over the past twenty-five years highlighting the advantages of doing so [23], [31]). Recently, however, the amount of research at the intersection of these fields has begun to dramatically increase [51], [52]. This can be primarily attributed to recent advances in technologies in this space that has not only significantly improved previously conceptualized concepts but also opened up newer intriguing avenues of human-robot interactions through the medium of virtual or augmented reality. Furthermore, these new works leverage AR and VR technologies in a variety of different ways, often in ways that blurs the boundaries between these two technologies. For example, VR technologies are being used not only to provide a view into simulated worlds, but also to provide alternate or augmented views of real-world environments, a domain typically the purview of AR technologies. Similarly, AR technologies are being used not only to annotate users' views of reality but also to explore simulation with virtual robots, a domain typically the purview of VR technologies.

This blurred boundary means that advances in VR/AR-for-HRI cannot be neatly categorized purely based on the technology used. Instead, we argue in this paper that such advances should be categorized in two primary ways.

First, we argue that the the most important dimension for categorizing advances in VR/AR-for-HRI is the way in which they present new opportunities for interactivity. Here, we are primarily interested in improvements to interactivity stemming from the use of different *interaction design elements*, such as interface elements and communication channels. Taking inspiration from the *Model-View-Controller* design pattern [29], we begin by arguing that interaction design elements (both in VR/AR interfaces as well as in HRI in general) can be categorized based on whether they improve (1) the expressivity of a user's view of robot's state; and/or (2) flexibility of the user's control of the robot, both of which are mediated by the complexity and depth of the robot's internal model. These serve as orthogonal dimensions on the space of interactivity, which we term the *Plane of Interaction*. We then present a taxonomy for conceptualizing *mixed-reality* interaction design elements in particular along this dimension.

Next, we argue that the second most important dimension for categorizing advances in VR/AR-for-HRI is the environment (e.g., reality vs. virtuality) over- or into-which mixed-reality interaction design elements are rendered. Due to the blurred boundary between the use of VR/AR technologies, this dichotomy is correlated with but ultimately distinct from the VR-vs-AR dichotomy. By combining the reality-virtuality continuum with the Plane of Interaction, we present a comprehensive framework for conceptualizing work within the field of VR/AR-for-HRI: we refer to this as the *Reality-Virtuality Cube of Interaction*.

Finally, we leverage this comprehensive framework to analyze previous work in VR/AR-for-HRI. Specifically, we examine as an initial case study the set of advances presented at VAM-HRI 2018[1], the 1st International Workshop on Virtual, Augmented, and Mixed-Reality for Human-Robot Interaction, and use this analysis to highlight recent research trends in the nascent VAM-HRI community.

The rest of the paper proceeds as follows. In Section II, we present the Plane of Interaction. In Sections III and IV, we introduce the concept of Interaction Design Elements and our taxonomy of Mixed-Reality Interaction Design Elements. In Sections V an VI we describe Milgram's Reality-Virtuality Continuum and how it combines with the Plane of Interaction to create the Reality-Virtuality Interaction Cube. In Section VII

[1]http://vam-hri.xyz/

we apply our framework to papers presented at VAM-HRI 2018. Finally, in Section VIII we conclude with potential directions for future work within the VAM-HRI community.

II. The Plane of Interaction

To reason about the ways in which interactive technologies can be advanced, we leverage the **Model-View-Controller** design pattern. This design pattern separates interactive systems into three pieces: the internal *model* of the interactive system, the user's *view* into that model, and the user's *controller* for effecting changes to that model [29].

This design pattern can easily be applied to interactive robots. Every robot has some internal state that can include matters-of-fact such as the robot's pose and battery level as well as cognitive constructs such as beliefs, goals, and intentions [20]. Moreover, robots that are truly interactive present opportunities for *view* into that internal model, and opportunities for some degree of *control* over that internal model. In particular, we argue that robots' potential for interaction depends on the *expressivity* of view into the robot's model (i.e., the scope of means used by the robot to passively or actively communicate its internal state), and the *flexibility* of control (i.e., the scope of means available for users to modify the robot's internal state).

Accordingly, a robot's level of interactivity can be conceptualized as a point on the *Plane of Interaction*, with two axes: *expressivity of view* (hereafter EV) and *flexibility of controller* (hereafter FC), as illustrated in Figure 1. Here, the focus on expressivity and flexibility is crucial. Simply logging additional low-level data without providing higher-level features or summaries, or new ways of viewing that data, would yield a limited gain in expressivity. Similarly, simply enabling complete, direct joysticking without providing opportunities to influence the robot's higher-level beliefs, desires, and intentions, would yield a limited gain in flexibility.

For example, as shown in the Figure 1, a robot employing manipulable visual cues of artifacts from its plan [15] allows for greater flexibility of control (e.g. the user can manipulate parts of the plan and initiate replanning on the part of the robot) as well as greater expressivity of view (e.g. beyond constructs confined to motion planning such as trajectories and areas) than just visually projecting its areas of influence onto the environment in mixed reality. More such examples, and their relation to the plane of interaction, are discussed in more detail later in Section VII, using instances from existing literature on human-robot interactions in mixed reality.

This plane does not explicitly include the robot's model. Instead, the complexity of the robot's model is implicitly represented by the scale of the plane. The more sophisticated a robot's model, the greater potential for EV and FC. That is, improving the complexity of a robot's model of the world (e.g., through the use of rich ontologies and Knowledge Processing frameworks [44]) does not *directly* improve interactivity, but instead increases the potential for improvement that can be gleaned from improvements to expressivity of EV and FC.

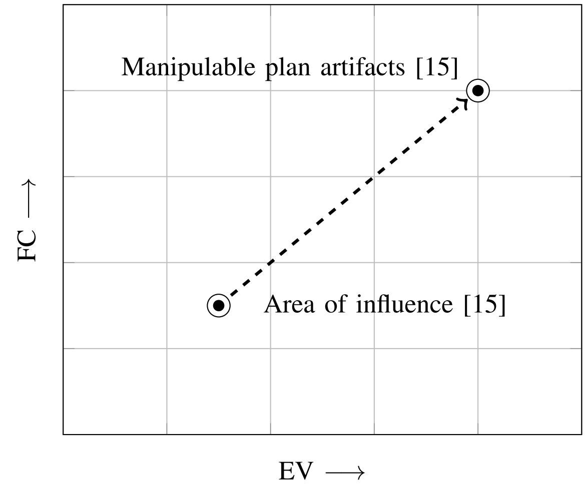

Fig. 1: The Plane of Interaction illustrating how the expressivity of view and flexibility of the controller vary based on the choice of mixed reality interaction design elements. We will discuss these in more detail in Section VII. Note that the positioning and scale of the points are merely illustrative and not to be taken literally.

III. Interaction Design Elements

We define *interaction design elements* as those components of a robot's design[2] that can be said to impact its interactivity and thus, its position on the plane of interaction. While it is likely infeasible to explicitly determine the position of a technology on this plane, it is nevertheless instructive to consider the formal relationship between interaction design elements and the position of a technology on this plane.

We define the impact of each design element on the robot's interactivity as $M \begin{bmatrix} \Delta_{EV} \\ \Delta_{FC} \end{bmatrix}$, where Δ_{EV} is the impact a design element has on the expressivity of the user's view into the robot's model, and Δ_{FC} is the impact a design element has on the flexibility of the user's control of the robot's model, both of which are scaled by M, a measure of complexity of the robot's internal model.

IV. Mixed Reality Interaction Design Elements

We term the subset of interaction design elements employed with virtual and augmented reality technologies as *mixed-reality* interaction design elements (MRIDEs). In this paper, we consider three principle categories of MRIDEs. To illustrate these categories, we will discuss examples presented together by [16] in a paper at VAM-HRI 2018.

User-Anchored Interface Elements: interface elements similar to those seen in traditional GUIs, anchored to points in the user's camera's coordinate system, and which do not move as the user changes their field of view. As an illustrative example, [16] present User-Anchored Interface Elements in the form of

[2]While this framework can be applied to any interactive technology, in this work we employ it specifically to interactive robots.

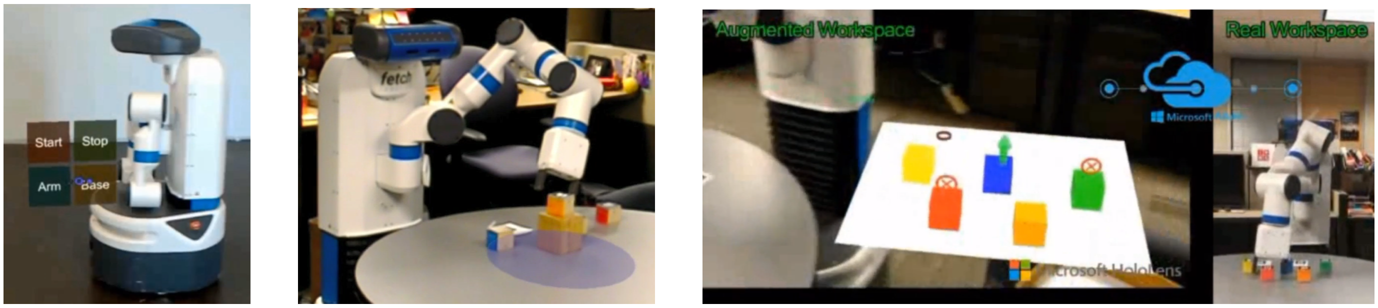

(a) A holographic control panel for the robot.

(b) Safety cues externalizing internal state information of the robot.

(c) Manipulable virtual objects used to specify intended actions and receive feedback from the human in the loop.

Fig. 2: Examples of interaction design elements [16] for human-robot interactions in mixed-reality. Used with permission.

an interactive control panel (Figure 2a). This enhances control of the robot, but not view into its internal state.

Environment-Anchored Interface Elements: interface elements anchored to points in the coordinate system of a robot or some other element of the environment, rather than anchored to the interface itself. As an illustrative example, [16] present visualizations anchored to a robot in order to display the area in which the robot is capable of moving. This enhances view of the robot, but not view.

Virtual Artifacts: 3D objects that can be manipulated by either humans or robots (or which may move under their own ostensible volition), or which may impact the behaviors of robots. Crucially, virtual artifacts in this category are recognizable (by humans) as additions to the underlying environment, but may not be recognizable as such to robots. For example, a robot may render an arrow into the environment which a human may then manipulate to change the robot's intended direction; an arrow recognizable by both parties as not part of the actual environment. In contrast, a human may render a virtual wall into the environment to restrict a robot's path, but this wall may or may not be identifiable as virtual to that robot. As an illustrative example, [16] present 3D virtual objects that serve two purposes. First, these virtual objects both reflect the robot's perception of the state of the world, thus enhancing view into the robot's internal state. Second, the user is able to manipulate the virtual objects in order to specify actions to be performed on their real-world analogues, or initiate replanning in the case of a more autonomous robot, thus enhancing control over its actions.

V. Reality-Virtuality Interface Continuum

Thus far we have described a taxonomy of interaction design elements and described how instances of each taxonomic class have been used by VAM-HRI researchers to enhance the expressivity of view and/or flexibility of control of a robot. In this section, we will consider the environment into which these elements are rendered. Specifically we consider two discrete types of environments, each falling at a different point along Milgram's Reality-Virtuality Continuum (Figure 3):

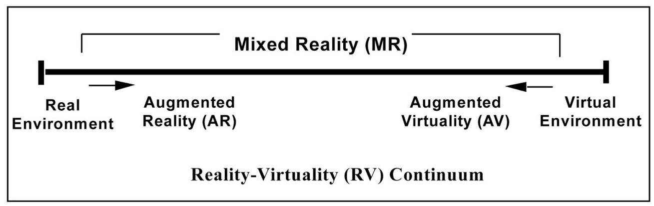

Fig. 3: Milgram's Reality-Virtuality Continuum

Reality: When using Augmented Reality technologies, such as Augmented Reality Head-Mounted Displays (AR-HMDs), projectors, or smart windshields, design elements are overlaid directly onto a user's local reality. This is also possible, however, using Virtual Reality technologies, by passing sensor data directly into a user's VR-HMD. Depending on the positioning of the sensor relative to to the display, this may either closely mimic AR, or serve as a window into a remote reality.

Virtuality: Finally, design elements may be overlaid onto completely virtual environments. This is typically the case when VR is used, especially when training or simulating robots before they are moved to the real world.

VI. The Reality-Virtuality Cube of Interaction

In this paper we have argued for two primary means of categorizing advancements in VR/AR-for-HRI: (1) the types of interaction design elements used and their associated benefits with respect to the Plane of Interaction, and (2) the underlying environments in which those interaction design elements are employed, e.g.,, reality or virtuality. Combining the 2D Plane of Interaction with the 1D Reality-Virtuality Continuum produces a 3-D space of interactive mixed-reality technologies, which we term the *Reality-Virtuality Cube of Interaction* (Figure 4). As with the Plane of Interaction, it is infeasible to try to identify the precise location within this 3-D space for a given interactive technology. However, this 3-D space can be leveraged as a useful tool for categorizing recently presented advancements in VR/AR-for-HRI.

For example, in Figure 4, we illustrate how a holographic control panel in AR [15] provides little flexibility and control

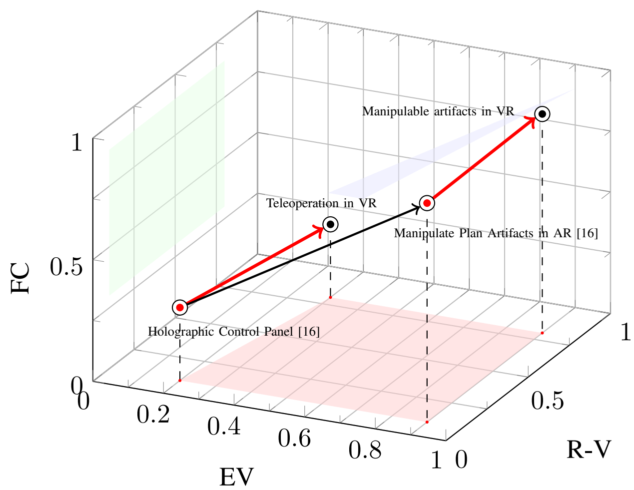

Fig. 4: The Reality-Virtuality Cube of Interaction illustrating how mixed reality interaction design elements extend the plane of interaction along the dimension of reality and virtuality. Again, the positioning and scale of the points are merely illustrative and not to be taken literally.

and expressivity of view, while a teleoperation interface in VR would be placed higher on the EV dimension by enhancing the view of the user. Similarly, plan artifacts in mixed reality, as we discussed before in Figure 1, situates us higher up in both FC and EV dimensions, across the spectrum of variations in the medium of AR and VR along the R-V dimension.

To demonstrate this utility, in the next section we use the Reality-Virtuality Cube of Interaction as a framework for analyzing the set of papers presented at VAM-HRI 2018.

VII. Survey: VAM-HRI 2018

To demonstrate the application of our proposed framework, we analyze the papers presented at the VAM-HRI workshop[3] at HRI 2018 [51], [52]. Using the proposed framework, we see that approximately two-thirds of papers leveraged VR or AR to enhance interactions in reality, while third of the papers leveraged VR or AR to enhance interactions in virtuality, as shown in Table I. Below, we will first discuss reality-enhancing applications, and then discuss virtuality-enhancing applications. Within each section, we will first discuss approaches that do *not* leverage any MRIDEs to enable these enhancements, and then discuss approaches that *do* leverage MRIDEs. Within each of these subsections, we will first discuss approaches that enhance neither view nor control, then approaches that enhance *either* view or control, and finally discuss approaches that enhance *both* view and control. For each approach leveraging MRIDEs, we will discuss the specific MRIDE used. In Table II, we further analyze the presented aproaches by dividing the presented MRIDEs into five categories: control, general information, spatial information, object information, and robot visualization.

[3]VAM-HRI Series: http://vam-hri.xyz/

	No MRIDEs used			MRIDEs used		
	∅	C+	V+	C+	V+	V+C+
Reality	0	0	7	3	10	4
Virtuality	3	3	0	0	4	0

TABLE I: VAM-HRI'18 papers, categorized by (1) whether Mixed Reality Interaction Design Elements (MRIDEs) are overlaid over the environment, (2) whether that environment is real or virtual, and (3) whether they improve flexibility of control (C+), expressivity of view (V+), both (V+C+), or neither (∅). A small number of papers were counted multiple times if they included multiple types of MRIDEs.

A. Reality-Enhancing Applications

1) No MRIDEs Used: Of the approaches enhancing interactions in reality, seven enhanced view without the use of mixed reality design elements. These all used a VR HMD for remote teleoperation, without displaying any additional graphics within the HMD. The VR HMD necessarily increased view (through direct visualization of remote robots' sensor data). We note that while the increases in expressivity of view enabled by these approaches *indirectly* enabled increases in flexibility of control as well (by allowing for remote teleoperation), we only consider the *direct* increases to view made by these approaches.

Rosen et al. present *ROS-Reality*, a ROS-based approach to robot teleoperation using the HTC Vive and its' associated controller [38], [47]. Bennett et al. leverage VR within a full-torso exosuit that allows for direct manipulation of a remote robot's head and arm motions for expressive teleoperation [10], [11]. Zhang et al. pair VR with a 360° camera to enable flexible viewing of remote environments without requiring rotation of the VR-HMD to correspond with rotation of physical elements on the robot [54], [55]. Oh et al., similarly pair VR with a 360° camera to enable flexible viewing of remote environments without requiring rotation of the VR-HMD to correspond with rotation of physical elements on the robot [33]. Tran et al. pair phone-based VR with the LeapMotion to enable hands-free, low-cost robot teleoperation [45]. Allspaw et al. present a robot teleoperation interface in which teleoperators view either a portrayal of reality reconstructed from rich point-cloud data and/or sets of VR-interactable camera feeds [1], [2]. Gaurav et al. also present a teleoperation interface based on point-cloud data, but instead focus on learning correspondences between the movements of the teleoperator within this interface and those of the teleoperated robot [19].

2) MRIDEs Used: Of the approaches enhancing interactions in reality, fourteen enhanced interactions using MRIDEs.

Control-Enhancing MRIDEs: Three approaches used control-enhancing augmentations, displaying environment- or user-anchored interface elements for controlling robots during VR/AR teleoperation. Arévalo-Arboleda et al. propose an approach towards teleoperating a local robot through Augmented Reality, where the user's view of the robot is augmented with control panels anchored to the robot and to interactable elements in the environment [5]. Oh et al. enhance flexibility of control by augmenting the view of a remote environment with a user-anchored control panel allowing users to specify their intended destination [33]. Finally, Chakraborti et al. enhances control of the robot by providing a control panel that is attached to the robot in the augmented view, allowing the user to start and stop execution of the robot's plan in progress or even seize control of the base or arm of the robot for more detailed control. [16].

View-Enhancing MRIDEs: Ten approaches used view-enhancing augmentations, including passive displays of robot trajectories or sensor data, active communicative displays, or even virtual robots in entirety.

Williams proposes the use of environment-anchored visualizations such as circles and arrows as stand-ins for traditional physical gestural cues within an active communication framework, as well as the use of simulated robot arms also for the purposes of robotic gestures [49], [50], [53]. Peters et al. present an approach to using not only simulated limbs, but entire robots simulated in AR, in order to study human-robot proxemics [6], [18], [34]. Katzakis et al. also present an approach toward visualizing entire simulated humanoid robots, anchored to real ground robots in the environment, and demonstrate a prototype of this approach within a simulated virtual environment [28]. Chakraborti et al. [14] takes the idea of enhancing interactions in humans and robots to vehicle to vehicle (v2v) communications and highlights how information can be effectively relayed from autonomous vehicles to nearby human drivers without causing information overload. Puljiz et al. present several ways of leveraging AR within autonomous warehouses [35], including environment-anchored control panels, maps, safety boundaries, and virtual objects used to highlight intended and/or occluded objects. Ben Amor et al. present an approach in which environment-anchored visualizations displayed via projector are used to convey task-related cues, including safety boundaries, task instructions and feedback, and highlighting of intended and/or occluded objects [3], [4], [18]. Chakraborti et al. [16] also provides task-related visual cues in the form of annotations of areas of influence, objects to be manipulated, awareness of peripheral and hidden objects, etc. and formalizes the notion of task-related cues in a *domain-independent* framework for visualizing artifacts of task plans (e.g. actions, states, intentions, etc.). They also provide a unique role of design elements in what they referred to as "projection-aware planning" where the elements are used not only for visualization of execution but also at the time of planning to generate plans that are easier to visualize. Zu Borgsen et al. visualize robots' trajectories and status information such as battery levels, as well as simulated robot parts like robot heads, as environment-anchored visualizations within AR-HMDs [57]. They also demonstrate an approach involving entire robots simulated in a VR CAVE environment. Cheli et al. present an approach towards using AR in middle-school robotics education, using a handheld tablet to visualize environment-anchored robot status information [17]. They also propose environment-anchored visualization of trajectories, boxes around detected objects, sensor readings, and occupancy grids, within AR. Bagchi et al. propose using AR to display robot-generated text, highlight objects, and present task feedback in noisy environments [9].

Control-And-View-Enhancing MRIDEs: Finally, four approaches used augmentations simultaneously enhancing both view and control. Two used AR UI elements to control and calibrate a virtual robot, and two provided virtual objects that could be interacted with to affect robot behavior.

Bagchi and Marvel present an approach towards calibrating robots, in which a tablet is used to visualize an environment-anchored simulated robot controlled with user-anchored interface items [9]. Chakraborti et al. [16] deploys manipulable virtual objects to annotate artifacts of a long term task plans thereby allowing the user to interact with them and either request for specific actions or initiate real-time replanning on the part of the robot (depending on the level on autonomy of the agent) during online plan execution based on the feedback from the user. Quintero et al. present an approach to visualizing trajectories in AR in a way such that those trajectory visualizations can be interacted with, and show how simulated robot arms can be visualized following those manipulated trajectories [36], [37]. Schönheits present an approach in which virtual robots are displayed over the physical world to make clear whether their positions are accurately calibrated, and user-anchored control panels that can be used to recalibrate the positions of those robots [39].

B. Virtuality-Enhancing Applications

1) No MRIDEs Used: Of the approaches enhancing interactions in virtuality, six did so without using MRIDEs, simply using VR as a window into a virtual environment. These six approaches focused either on using VR to allow humans to train robots or on using VR to train humans to interact with robots or study perceptions of virtual robots.

Three of these increased control by enabling humans to control virtual robots. Both Whitney et al. and Stramandinoli et al. present virtual reality training environments where humans can train robots by demonstration before the learned policies are executed in the real world [43], [48]. And Sportillo et al. present a virtual environment for training humans to interact with autonomous vehicles [41], [42].

The other three approaches in this category allowed observation of uncontrollable virtual robots. Hansen et al. present a virtual reality training environment for the meat processing industry in which humans can observe the behavior of co-located industrial robots [24]. Goedicke et al. present a virtual reality training environment in which users view themselves within autonomous cars [21], [22]. And Iuzzolino et al. present a photorealistic environment in which robots learn to navigate before moving to similar real-world environments [26], [27].

2) MRIDEs Used: Finally, four approaches enhanced expressivity of view in virtuality. These primarily involved large-scale maritime or aviation contexts in which it was more helpful to see a top-down view of the larger maritime or

Category	MRIDE	User Anchored	Environment Anchored	Manipulable Virtual Object
Control	Control panels	[33], [9], [39]	[5]*, [35]	
General Information	Robot-Generated Text		[9]*	
	Status Information		[42]*, [57], [17]	
	Orientation Information		[30]	
	Task Instructions / Feedback		[9]*, [42]*, [3]	
Spatial Information	Trajectories		[17]*, [32], [57], [30]	[37]
	Sensor Readings		[17]*	
	Maps and Region highlighting		[35], [30], [25]	
	Occupancy Grids		[17]*	
	Safety Boundaries		[35], [3]	
Object Information	Circles, Boxes, and Arrows		[49]*, [17]*	
	Far-off object view enhancement		[30]	
	Intended object highlighting		[9]*, [3], [35]	
	Occluded object display		[35], [3]	
Robot Visualization	Virtual robots		[28], [9], [39]	[57], [34]
	Virtual appendages		[49]*, [57], [37]	

TABLE II: MRIDEs presented or proposed (*denotes proposed rather than presented approaches) at VAM-HRI 2018).

aerial region than just the perspective of the single unmanned surface vehicle or drone, and in which helpful information was overlaid on the canvas of the open maritime or air space. Novitzky et al. present a VR environment for studying human teaming with semi-autonomous robotic boats, in which they use simple environment-anchored annotations such as circles delineating goal regions and lines indicating trajectories [32]. Lager et al. present a VR environment for remote operation of unmanned surface vehicles, in which a wide variety of environment-anchored annotations are displayed, including geographical data (i.e., sea charts), directional information, trajectories, and blown-up imagery of far-off objects [30]. Haring et al. present a VR environment for controlling swarms of aerial drones, in which are displayed environment-anchored annotations delineating important spatial regions [25]. Finally, Sportillo et al. propose the use of AR-like visualizations rendered within their VR training simulator to enable better user training, specifically environment-grounded task instructions and status information [41], [42].

VIII. Conclusion

In this paper, we have presented a novel framework for characterizing interactive technologies based on the Model-View-Controller paradigm, and demonstrated how it can be combined with the reality-virtuality continuum to create a 3D-space of interactive technologies that leverage Augmented and Virtual Reality. As a proof of concept demonstration of the utility of this framework, we analyzed the papers presented at VAM-HRI 2018. This analysis suggests that most previous approaches within VR/HR-HRI have sought to improve either view *or* control, whereas many new opportunities may exist in the development of approaches that improve both view *and* control. Moreover, we have presented an initial taxonomy of Mixed Reality Interaction Design Elements presented in papers at VAM-HRI 2018. This initial taxonomy shows that most previous work has focused on presentation of environment-anchored annotations, while much fertile ground on the use of manipulable virtual objects remains open for exploration. In future work we hope to use the analysis in this paper as the starting point for a survey of the field of VR/AR-for-HRI. Moreover, we hope that the novel concepts presented in this work will serve as useful frameworks for future VR/AR-for-HRI researchers hoping to categorize their own work relative to other recent work in the field.

References

[1] Jordan Allspaw, Jonathan Roche, Nicholas Lemiesz, Michael Yannuzzi, and Holly A Yanco. Remotely teleoperating a humanoid robot to perform fine motor tasks with virtual reality–. In *Proceedings of the 2018 Conference on Waste Management*, 2018.

[2] Jordan Allspaw, Jonathan Roche, Adam Norton, and Holly A Yanco. Remotely teleoperating a humanoid robot to perform fine motor tasks with virtual reality–. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[3] Heni Ben Amor, Ramsundar Kalpagam Ganesan, Yash Rathore, and Heather Ross. Intention projection for human-robot collaboration with mixed reality cues. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[4] Rasmus S Andersen, Ole Madsen, Thomas B Moeslund, and Heni Ben Amor. Projecting robot intentions into human environments. In *Proceedings of the 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 294–301. IEEE, 2016.

[5] Stephanie Are'valo-Arboleda, Max Pascher, and Jens Gerken. Opportunities and challenges in mixed-reality for an inclusive human-robot collaboration environment. In *The 1st International Workshop on Virtual, Augmented, and Mixed Reality for Human-Robot Interactions (VAM-HRI)*, pages 83–86, Chicago, USA, 2018.

[6] Vanya Avramova, Fangkai Yang, Chengjie Li, Christopher Peters, and Gabriel Skantze. A virtual poster presenter using mixed reality. In *International Conference on Intelligent Virtual Agents*, pages 25–28. Springer, 2017.

[7] Ronald Azuma, Yohan Baillot, Reinhold Behringer, Steven Feiner, Simon Julier, and Blair MacIntyre. Recent advances in augmented reality. Technical report, Naval Research Lab, Washington, DC, 2001.

[8] Ronald T Azuma. A survey of augmented reality. *Presence: Teleoperators & Virtual Environments*, 6(4):355–385, 1997.

[9] Shelly Bagchi and Jeremy A Marvel. Towards augmented reality interfaces for human-robot interaction in manufacturing environments. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[10] Maxwell Bennett, Tom Williams, Daria Thames, , and Matthias Scheutz. Differences in interaction patterns and perception for teleoperated and autonomous humanoid robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017.

[11] Maxwell Bennett, Tom Williams, Daria Thames, and Matthias Scheutz. Investigating interactions with teleoperated and autonomous humanoids using a suit-based vr teleoperation interface. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI*, 2018.

[12] Mark Billinghurst, Adrian Clark, and Gun Lee. A survey of augmented reality. *Foundations and Trends in Human–Computer Interaction*, 8(2-3):73–272, 2015.

[13] C Burdea Grigore and P Coiffet. *Virtual reality technology*. London: Wiley-Interscience, 1994.

[14] Tathagata Chakraborti, Andrew Dudley, and Subbarao Kambhampati. v2v communication for augmenting reality enabled smart huds to increase situational awareness of drivers. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[15] Tathagata Chakraborti, Sarath Sreedharan, Anagha Kulkarni, and Subbarao Kambhampati. Projection-aware task planning and execution for human-in-the-loop operation of robots in a mixed-reality workspace. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4476–4482. IEEE, 2018.

[16] Tathagata Chakraborti, Sarath Sreedharan, Anagha Kulkarni, and Subbarao Kambhampati. Projection-aware task planning and execution for human-in-the-loop operation of robots in a mixed-reality workspace. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018. Also appeared in the Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI).

[17] Mark Cheli, Jivko Sinapov, Ethan E. Danahy, and Chris Rogers. Towards an augmented reality framework for k-12 robotics education. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[18] Ramsundar Kalpagam Ganesan, Yash K Rathore, Heather M Ross, and Heni Ben Amor. Better teaming through visual cues. *IEEE Robotics & Automation Magazine*, 2018.

[19] Sanket Gaurav, Zainab Al-Qurashi, Amey Barapatre, and Brian Ziebart. Enabling effective robotic teleoperation using virtual reality and correspondence learning via neural network. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[20] Michael Georgeff, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. The belief-desire-intention model of agency. In *International Workshop on Agent Theories, Architectures, and Languages*, pages 1–10. Springer, 1998.

[21] David Goedicke, Jamy Li, Vanessa Evers, and Wendy Ju. Vr-oom: Virtual reality on-road driving simulation. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[22] David Goedicke, Jamy Li, Vanessa Evers, and Wendy Ju. Vr-oom: Virtual reality on-road driving simulation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 165. ACM, 2018.

[23] Scott A Green, Mark Billinghurst, XiaoQi Chen, and J Geoffrey Chase. Human-robot collaboration: A literature review and augmented reality approach in design. *International Journal of Advanced Robotic Systems*, 5(1):1, 2008.

[24] Laerke I. N. Hansen, Niklas Vinther, Lukas Stranovsky, Mark P. Philipsen, Haiyan Wu, and Thomas B. Moeslund. Collaborative meat processing in virtual reality. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[25] Kerstin S. Haring, Victor Finomore, Dylan Muramato, Nathan L. Tenhundfeld, Mormon Redd, James Wen, and Brian Tidball. Analysis of using virtual reality (vr) for command and control applications of multi-robot systems. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[26] Michael L Iuzzolino, Michael E Walker, and Daniel Szafir. Virtual-to-real-world transfer learning for robots on wilderness trails. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 576–582. IEEE, 2018.

[27] Michel L. Iuzzolino, Michael E. Walker, and Daniel Szafir. Virtual-to-real-world transfer learning for robot navigation. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[28] Nicholas Katzakis and Frank Steinicke. Excuse me! perception of abrupt direction changes using body cues and paths on mixed reality avatars. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 147–148. ACM, 2018.

[29] Glenn E Krasner, Stephen T Pope, et al. A description of the model-view-controller user interface paradigm in the smalltalk-80 system. *Journal of object oriented programming*, 1(3):26–49, 1988.

[30] Mårten Lager, Elin A. Topp, and Jacek Malec. Remote operation of unmanned surface vessel through virtual reality - a low cognitive load approach. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[31] Paul Milgram, Shumin Zhai, David Drascic, and Julius Grodski. Applications of augmented reality for human-robot communication. In *Intelligent Robots and Systems' 93, IROS'93. Proceedings of the 1993 IEEE/RSJ International Conference on*, volume 3, pages 1467–1472. IEEE, 1993.

[32] Michael Novitzky, Michael R. Benjamin, Paul Robinette, Hugh R.R. Dougherty, Caileigh Fitzgerald, and Henrik Schmidt. Virtual reality for immersive simulated experiments of human-robot interactions in the marine environment. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[33] Yeonju Oh, Ramviyas Parasuraman, Tim McGraw, and Byung-Cheol Min. 360 vr based robot teleoperation interface for virtual tour. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[34] Christopher Peters, Fangkai Yang, Himangshu Saikia, Chengjie Li, and Gabriel Skantze. Towards the use of mixed reality for hri design via virtual robots. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[35] David Puljiz, Gleb Gorbachev, and Björn Hein. Implementation of augmented reality in autonomous warehouses: Challenges and opportunities. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[36] Camilo Perez Quintero, Sarah Li, Matthew KXJ Pan, Wesley P Chan, HF Machiel Van der Loos, and Elizabeth Croft. Robot programming through augmented trajectories in augmented reality. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1838–1844. IEEE, 2018.

[37] Camilo Perez Quintero, Sarah Li, Cole Shing, Wesley Chan, Sara Sheikholeslami, H.F. Machiel Van der Loos, and Elizabeth Croft. Robot programming through augmented trajectories. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[38] Eric Rosen, David Whitney, Elizabeth Phillips, Daniel Ullman, and Stefanie Tellex. Testing robot teleoperation using a virtual reality interface with ros reality. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[39] Manfred Schönheits and Florian Krebs. Embedding ar in industrial hri applications. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[40] William R Sherman and Alan B Craig. *Understanding virtual reality: Interface, application, and design*. Morgan Kaufmann, 2018.

[41] Daniele Sportillo, Alexis Paljic, Mehdi Boukhris, Philippe Fuchs, Luciano Ojeda, and Vincent Roussarie. An immersive virtual reality system for semi-autonomous driving simulation: a comparison between realistic and 6-dof controller-based interaction. In *Proceedings of the 9th International Conference on Computer and Automation Engineering*, pages 6–10. ACM, 2017.

[42] Daniele Sportillo, Alexis Paljic, Luciano Ojeda, Giacomo Partipilo, Philippe Fuchs, and Vincent Roussarie. Training semi-autonomous vehicle drivers with extended reality. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[43] Francesca Stramandinoli, Kin Gwn Lore, Jeffrey R. Peters, Paul C. O'Neill, Binu M. Nair, Richa Varma, Julian C. Ryde, Jay T. Miller, and Kishore K. Reddy. Robot learning from human demonstration in virtual reality. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[44] Moritz Tenorth and Michael Beetz. KnowrobâĂŤknowledge processing for autonomous personal robots. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 4261–4266. IEEE, 2009.

[45] Nhan Tran, Josh Rands, , and Tom Williams. A hands-free virtual-reality teleoperation interface for wizard-of-oz control. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.

[46] DWF Van Krevelen and Ronald Poelman. A survey of augmented reality technologies, applications and limitations. *International journal of virtual reality*, 9(2):1, 2010.
[47] David Whitney, Eric Rosen, Elizabeth Phillips, George Konidaris, and Stefanie Tellex. Comparing robot grasping teleoperation across desktop and virtual reality with ros reality. In *Proceedings of the International Symposium on Robotics Research*, 2017.
[48] David Whitney, Eric Rosen, and Stefanie Tellex. Learning from crowdsourced virtual reality demonstrations. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.
[49] Tom Williams. A framework for robot-generated mixed-reality deixis. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.
[50] Tom Williams, Matthew Bussing, Sebastian Cabrol, Elizabeth Boyle, and Nhan Tran. Mixed reality deictic gesture for multi-modal robot communication. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction*, 2019.
[51] Tom Williams, Daniel Szafir, Tathagata Chakraborti, and Heni Ben Amor. Report on the 1st international workshop on virtual, augmented, and mixed reality for human-robot interaction (VAM-HRI). *AI Magazine*, 2018.
[52] Tom Williams, Daniel Szafir, Tathagata Chakraborti, and Heni Ben Amor. Virtual, augmented, and mixed reality for human-robot interaction. In *Companion of HRI 2018*, pages 403–404. ACM, 2018.
[53] Tom Williams, Nhan Tran, Josh Rands, and Neil T. Dantam. Augmented, mixed, and virtual reality enabling of robot deixis. In *Proceedings of the 10th International Conference on Virtual, Augmented, and Mixed Reality (VAMR)*, 2018.
[54] Jingxin Zhang, Eike Langbehn, Dennis Krupke, Nicholas Katzakis, and Frank Steinicke. A 360âŮę video-based robot platform for telepresent redirected walking. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.
[55] Jingxin Zhang, Eike Langbehn, Dennis Krupke, Nicholas Katzakis, and Frank Steinicke. Detection thresholds for rotation and translation gains in 360Âř video-based telepresence systems. *IEEE transactions on visualization and computer graphics*, 24(4):1671–1680, 2018.
[56] Feng Zhou, Henry Been-Lirn Duh, and Mark Billinghurst. Trends in augmented reality tracking, interaction and display: A review of ten years of ismar. In *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 193–202. IEEE Computer Society, 2008.
[57] Sebastian Meyer zu Borgsen, Patrick Renner, Florian Lier, Thies Pfeiffer, and Sven Wachsmuth. Improving human-robot handover research by mixed reality techniques. In *Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*, 2018.