


MSc. Thesis

Visual Feedback for
Haptic Assisted Teleoperation
of an Industrial Robot
With Dross Removal as a Use Case

Technische Universiteit Delft



Visual Feedback for Haptic Assisted Teleoperation of an Industrial Robot With Dross Removal as a Use Case

by

A. van den Berg

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday November 12, 2020 at 10:30.

Student number: 4228391
Master Track: BioMechanical Design
Project duration: November 1, 2019 – November 12, 2020
Thesis committee: Prof. dr. ir. D. A. Abbink, TU Delft, supervisor
Dr. ir. L. Peternel, TU Delft, supervisor
Ir. J. Hofland, Heemskerk Innovative Technology, supervisor
Dr. ir. C. J. M. Heemskerk, Heemskerk Innovative Technology, supervisor
Dr. ir. J. Kober, TU Delft, external member

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

The ability to perform tasks from a distance through teleoperation has the obvious benefit that humans can be removed from dangerous or remote environments. However, this benefit comes at the cost of our ability to directly interact with and perceive that environment. Through the use of haptic interfaces, great advances have been made to help to overcome this perceptual barrier by enabling us to *feel* these environments, without ever being present. Haptic Assistance (HA) has provided some additional magic to these interfaces, as virtual forces can be implemented to guide us in the process of teleoperation. During my internship at Heemskerk Innovative Technology I got the opportunity to design such a system. However, I found out this haptic guidance information is often difficult to interpret, especially for inexperienced operators. In an attempt to improve this understanding, literature was investigated to identify what visual aids have been implemented in shared control teleoperation systems (entire study is available at <http://repository.tudelft.nl/>). It was found that, although many visual cues have been implemented, they are rarely evaluated, and their implementation in haptic teleoperation is quite rare. In this MSc. thesis I set out to bridge some of the gaps in this literature, improving our understanding of the design of HA teleoperation interfaces and furthering our ability to perform tasks from a distance.

This thesis is submitted in partial fulfilment for the requirements of my Master's degree in Biomechanical Design at the Delft University of Technology. The research presented herein was conducted under the supervision of Ir. Jelle Hofland and Dr. Cock J. M. Heemskerk of Heemskerk Innovative Technology B.V., and Prof. dr. ir. David A. Abbink and Dr. ir. Luka Peternel of the Department of Cognitive Robotics at the Faculty of Mechanical, Maritime and Materials Engineering.

*A. van den Berg
Delft, November 2020*

Acknowledgements

First, I would like to thank my supervisors David Abbink, Luka Peternel, Jelle Hofland and Cock Heemskerk. I would like to thank David for his ever-present enthusiasm, critical advice and supernatural ability to somehow always see the bigger picture, Luka for his motivating words and for sharing his thoughts from a different perspective during the progress meetings, Jelle for his daily advice and efforts to always find the time to guide and assist me, and Cock for providing me with the opportunities, advice, and the awesome and unique work environment that is the Heemskerk Innovative Technology office. Moreover, I thank all my colleagues for their part in creating this work environment by means of support, advice, and friendship.

In addition, I owe a special thanks to my amazing friends, brother and sisters and of course my parents, for their unconditional love and support and the unforgettable experiences we have lived together. Finally, I am indescribably grateful to my girlfriend Judith, who has been by my side and supporting me for over ten years now. Your cups of coffee in the morning after a long night of work, your love, patience, and uncompromising trust in me have given me the confidence and motivation and that I needed to get to this point.

Contents

- 1 Scientific Paper 1
- A Unity Scene 15
 - A.1 Video Demonstration 15
 - A.2 Unity Project Download. 15
 - A.3 External Assets 15
 - A.4 Visual Features 16
 - A.4.1 Realistic looking materials 16
 - A.5 Robot Control. 17
 - A.5.1 Kinematics. 17
 - A.5.2 Haptic Control 18
 - A.5.3 Haptic Assistance 18
 - A.6 Dross 19
 - A.7 Logging and Data processing 20
- B Virtual Fixtures Shader 21
- C Participant Instructions 25
- D Participant Consent Form 32
- E Questionnaires 34
- F Shapiro Wilk Test of Normality Results 36
- G Two-way RM ANOVA Results 37
- H Other Results 38
- Bibliography 41

1

Scientific Paper

Visual Feedback for Haptic Assisted Teleoperation of an Industrial Robot

Alex van den Berg

Supervised by:

Jelle Hofland, Cock J.M. Heemskerk, Luka Peternel, and David A. Abbink

Abstract—Haptic assistance (HA) has been shown to be useful in a wide variety of applications by providing the benefits of automation while keeping the human in the loop. In recent years, providing additional visual feedback has been found to provide benefits to HA interfaces by complementing their advantages and mitigating their weaknesses. However, the implementation and understanding of this interaction are still quite limited. In this study, we provide new insights into this interaction by extending these findings to the HA teleoperation of an industrial robot, in which the HA is comprised of a set of Virtual Fixtures (VFs). Two methods of providing additional visual feedback are implemented for the use case of dross removal and are evaluated in a human factors experiment. The first method is the implementation of a set of visual cues, designed to complement the HA and force feedback. The second method is the use of a head-mounted display (HMD), instead of a desktop monitor, providing additional depth information and an increased sense of immersion. Both methods proved to be beneficial, but only in certain aspects of the operation. The visual cues were found to significantly improve safety in terms of peak collision force, whereas the HMD significantly improves the performance. Additional analysis suggests the use of an HMD causes improvements in the manipulability of the interface. Furthermore, improved scores in the van der Laan questionnaires and the user preference indicate an increased user acceptance due to the implementation of either of the methods. This study provides additional insight into the importance of visual feedback for HA and provides two methods to take advantage of its potential benefits in the teleoperation of an industrial robot.

Index Terms—Teleoperation, Haptic Assistance, Virtual Fixtures, Haptic Feedback, Visual Feedback, Dross Removal, VR, HMD



1 INTRODUCTION

DANGEROUS, remote or unreachable environments such as nuclear facilities and space stations still require dexterous tasks such as maintenance to be performed on-site. In these situations, direct manipulation can be undesired or simply not possible. Here, teleoperation can be a solution, as it allows humans to perform dexterous tasks in such environments while remaining in the loop. Full automation has been suggested as a solution, but is not (yet) capable enough to provide satisfactory results. Although fully autonomous robots can perform very well on repetitive tasks in standardized environments, when physically interacting with complex and varying environments their capabilities are limited [1]. If such a task would be automated, this can lead to problems such as misuse, disuse, and abuse, ultimately resulting in a decrease in satisfaction, performance, and safety [2] [3]. Keeping the human in the loop offers the ability to overcome these issues [4], as it allows operators to maintain their judgment, skill, attention, and ability to resolve unexpected situations [5] [6].

The major drawback that comes with teleoperation is that it is much more difficult than direct manipulation due to delays, limited sensory feedback, as well as a mismatch in view-points [7] [8]. This results in a limited performance, accuracy, and situational awareness [9] (which is critical for producing effectual robot behavior [10]). Supplying operators with haptic feedback from the environment has been shown to help in this regard, leading to improved task performance [11] [12] and reduced workload [13]. However, accurately rendering these forces often proves difficult due to technical reasons [14] and full transparency has not yet been achieved.

Haptic Assistance (HA) has been shown to be a promising compromise between manual teleoperation and full automation [15] [14]. In HA, guidance forces are provided to the operator using a haptic interface, with the goal of intuitively combining human intelligence and creativity with the benefits of automation systems [17]. One way in which HA can be supplied is in the form of virtual fixtures [18] (VFs). VFs are software generated force and position signals that can be either pushing the operator away from designated areas (forbidden region virtual fixtures, or FRVF) or guiding the operator along some desired path (guidance virtual fixtures, or GVF) [19]. HA has been successfully implemented in teleoperation to aid in protecting areas using FRVFs [19] [20] and in providing guidance to a certain reference position or path using GVFs [21] [22]. This has led to improvements such as increasing task performance and accuracy, and reducing mental workload.

- *A. van den Berg is with the Department of Biomechanical Engineering, Faculty 3mE; Delft University of technology, Mekelweg 2, 2628 CD Delft, The Netherlands, and with Heemskerk Innovative Technology B.V., Mijnbouwstraat 120, 2628 RX Delft, The Netherlands. E-mail: Alex_van_den_berg@hotmail.com*
- *C.J.M. Heemskerk and J. Hofland are with Heemskerk Innovative Technology B.V., Mijnbouwstraat 120, 2628 RX Delft, The Netherlands.*
- *L. Peternel and D.A. Abbink are with the Department of Cognitive Robotics, Faculty 3mE; Delft University of technology, Mekelweg 2, 2628 CD Delft, The Netherlands.*

1.1 Problem Statement

Despite the benefits brought by HA, part of the issues found in human-automation interaction remain unsolved [23]. One of these issues is inappropriate *trust* in the automation, potentially leading to inappropriate reliance. This in turn can cause problems such as misuse and disuse of the automation, as previously mentioned. Additionally, the implementation of HA introduces *conflicts* between the human and automation. These conflicts occur when the intention of the automation and the human are misaligned. Though not always the case, these conflicts are generally undesired [24] as they can lead to a deterioration in situational awareness and reliance [2] [25].

These issues can be summarized to arise from two separate problems: either the human does not understand the automation, or the automation does not understand the human [15]. This mismatch in understanding has been found in several HA interfaces, and is thought to be the cause of a low *user acceptance* [26] [20] [27]. As the degree of automation increases, it becomes critical that the operator has access to information about what the automated agents are doing and what they will be doing next [28]. Failure to provide this information may lead to operators experiencing difficulty in getting the automation to do what they want and have a poor understanding of how the automation works. However, the amount of information complexity that can be shared through haptic forces is limited [29] [30]. Consequently, as more information is being conveyed through these haptic forces, the information may become ambiguous and more difficult to interpret, potentially resulting in the problems stated above.

Only recently, researchers have started investigating how visual feedback can be used to provide operators with the information needed for solving these problems. The two main methods of how this is achieved are by either providing visual cues or by changing the type of display that is used.

The implementation of visual cues in HA interfaces has been found to result in improvements in user acceptance, safety, and performance in the aviation [31] [20], automotive [29] [32], and neurosurgery [27] domains. Although the use of visual cues in HA looks promising, their effects and implementation in HA teleoperation (especially that of an industrial robot) has, to the best of author's knowledge, not yet been investigated.

The HA in this study is comprised of a set of VFs, designed to guide and safeguard the operation only in key areas for the described use case of dross removal (Section 2.6). To complement these VFs (and other sources of haptic feedback) a set of visual cues is proposed and their effects are evaluated (Section 2.3). Their design has been based on previous work, where available, and will be further explained in Section 2.

Another promising method of providing operators with additional information using visual feedback is the implementation of a Head-Mounted Display (HMD). HMDs

provide more depth information [33], which is an important factor in teleoperation [9]. Additionally, HMDs allow for an increased sense of immersion, which improves teleoperation performance [34] [35].

The use of HMDs in teleoperation has gained popularity in the last few years, resulting in improvements in performance and user acceptance [36] [37] [38]. However, most of these methods have used the hand tracking of the virtual reality (VR) interface as an input method. As a result, these methods can not take advantage of (grounded) force feedback, impeding the implementation of HA. The effect of the use of an HMD in combination with HA is not trivial and is also, to the best of author's knowledge, not yet investigated. One possible reason for this is that using a haptic interface in combination with an HMD comes with the additional challenge that operators are not able to directly see the haptic interface they use.

This study proposes a way to overcome this challenge by supplying the operator with a virtual representation of this haptic interface. Using this method, the effects of the implementation of an HMD in HA teleoperation are evaluated.

1.2 Research Objectives

In this study, we propose two novel methods of improving visual feedback in HA teleoperation of an industrial robot using VFs. The first method is the implementation of a set of visual cues, specifically designed to complement the HA in the proposed use case (Section 2.3). The second method is the implementation of an HMD instead of a desktop monitor. Both methods, and the interaction between them, have been evaluated in a two-way human factors experiment, performed in a virtual environment. The HA is implemented in the form of VFs, based on previous research for similar interfaces.

Furthermore, with the evaluation of these two methods, we aim to provide new insights into the importance of visual feedback design in HA teleoperation. Lastly, some recommendations will be made, regarding their applicability for the proposed use case.

1.3 Hypotheses

Based on the related work, it was hypothesized that the implementation of the visual cues would help operators to improve in both *task performance* and *safety* of the operation, as well as to increase *user acceptance* (H1). Similarly, it was hypothesized that the use of an HMD over the desktop monitor would help to improve *task performance* and *user acceptance* (H2). Note that no safety improvements are expected, as there is no evidence found for this in previous research.

Furthermore, it was hypothesized that there is no interaction between these factors, such that these improvements are present, regardless of the state of the other factor (H3). The rationale behind this last hypothesis is that even though the information by the visual cues and the HMD partially overlap (both provide additional depth information), each of the methods has its own benefits. The main additional benefit of visual cues is that they provide information about the forces and the VFs, whereas the HMD provides a better sense of immersion.



Fig. 1. Pictures of the experiment setup. The participant is holding the stylus of the Touch haptic device and is wearing the HTC Vive HMD.

2 METHODS

2.1 Participants

Sixteen participants (4 female, 12 male) between 20 and 52 years old ($M = 24.8, SD = 7.6$) volunteered for the experiment. All participants gave their informed consent prior to the experiment. The setup and experiments were approved by the local ethics committee of the Delft University of Technology. Additionally, participants were asked how much experience they had with teleoperation. In response, 6 participants reported to have never done it, 2 have about 1 hour of experience, 6 have about 10 hours of experience, and 2 have about 1 day of the experience. Additionally, the participants were asked about their experience with video gaming. Most participants have more than 10 weeks of experience (12 participants), 2 participants reported to have about 10 weeks of experience and the last 2 participants reported to have about 10 hours of gaming experience.

2.2 Apparatus

The experiment was performed using the Touch Haptic Device by 3D Systems¹ (shown in Figure 1). This is a 6-DOF commanding haptic device that provides force feedback in 3-DOF (translations). The device has a workspace of 160 width \times 120 height \times 70 depth mm. It has a resolution of about 0.055 mm and can exert a maximum force of 3.3 N. The HMD device is an HTC Vive². It has a resolution of 1080 \times 1200 pixels per eye, a refresh rate of 90Hz, and a field of view of 110 degrees. The desktop monitor is a regular 23" desktop monitor with a resolution of 1920 \times 1080 pixels, and a 60Hz refresh rate. Both are shown in Figure 1.

The simulation was built and ran using Unity Game Engine (ver. 2019.3.9f1), on a desktop with a Intel(R) Core(TM) i7-8700 CPU @ 3.20GHz and an NVIDIA GeForce GTX 1660 Ti GPU, resulting in a refresh/sampling rate of about 100Hz.

2.3 Use Case

The task that has been chosen for this study is the removal of dross from the zinc bath in a continuous galvanizing line (CGL). Dross is a floating, solid contamination in the zinc bath that needs to be removed to maintain a high-quality coating. Its removal is a labor-intensive job with poor

work conditions, potential safety hazards, and difficulty in controlling operating costs and quality [39]. Furthermore, the environment is subject to changes, such as the liquid metals solidifying on the tools and environment and the fluctuation of liquid level in the zinc bath. Additionally, failing to adequately perform this dross removal operation could result in big financial losses.

These properties make it so that this task can potentially benefit from HA teleoperation, making it a suitable use case for this study.

2.4 Task Description

Participants were instructed to remove as much dross as they can within the time limit of five minutes while ensuring safe operation. An overview of the simulated environment is shown in Figure 2.

The dross is removed by controlling the industrial robot using the Touch Haptic Device. The control point is set to the wrist joint of the scoop (as if you are holding the scoop at the end of the rod). Using the scoop, the dross particles are removed from the zinc bath, and deposited into one of the two dross deposit bins.

The participants are instructed to cause as little disturbance to the bath as they can. This is important, as this disturbance leads to a deterioration of the quality of the coating applied in the galvanization process. Participants are informed this disturbance is minimized by minimizing the submerged volume of the scoop, and the velocity with which the scoop is moved through the bath. Furthermore, participants are instructed to try to avoid collisions and minimize the collision force when collisions do occur. Additionally, participants are instructed to move slowly and carefully and are enforced to do this in two ways. The first is that when moving quickly the dross particles "spill out" of the scoop, and the second is that fast movements result in high collision forces when a collision does occur (something which they are instructed to avoid). Lastly, two boundary conditions (BCs) are set, which will stop the current repetition if they are violated. These BCs are:

- *Maximum collision force.* In the real application, large collision forces might damage the robot and its surroundings. For this reason, a maximum collision force of 40 N is set. The value of this maximum is set in such a way that it is not easily exceeded by accident, as long as the robot is operated in a slow and controlled manner.
- *Collision with the steel strip.* If in a real CGL the robot would collide with the steel strip, it would damage the strip and potentially result in an shutdown of the entire galvanizing line, resulting in big financial losses. In the experiment, touching the steel strip stops the current repetition. The HA pushes the operator away from the steel strip. Nonetheless, they might be motivated to approach it, as there are dross particles in close proximity.

2.5 Design of Simulated Environment

The entire environment and input/output devices are set-up within the Unity game engine. Virtual models of

1. <https://www.3dsystems.com/haptics-devices/touch/specifications>

2. <https://developer.vive.com/resources/knowledgebase/vive-specs/>

the environment and robot are imported into the scene and are assigned physics colliders so that they can interact with each other. The robot model is a (ROS-Industrial) URDF of an ABB 6640³, imported using the ROS# plugin⁴. The graphics and layout were designed to be a realistic representation of the environment for dross removal, while keeping the rendering and computational costs to a minimum. Because of this, the simulation is able to run with about 100 frames per second, ensuring stability and comfort when using the HMD and haptic device.

The dross particles are implemented using the Obi Fluid⁵ Unity asset and tweaked to most accurately resemble actual dross behavior. There is a total of 750 dross particles placed in the bath. This number is chosen such that a skilled operator is able to remove about half the dross within the given time limit. This is important because it challenges operators to remove the dross from all areas in the bath, without running out of dross before the time limit is passed and thus preventing ceiling effects in the resulting data.

The *Touch* haptic device was integrated within Unity using a slightly modified version of the 3D Systems' official OpenHaptics Unity plugin⁶. This plugin communicates directly with the device drivers, maintaining high refresh rates. The plugin was modified such that upon pressing a button on the stylus of the haptic device, a connection is made between the haptic device's stylus and the wrist joint of the robot (see Figure 3), coupling their movements and rotations relative to the point at which this connection has been made. The (translational) workspace of the haptic device is scaled in such a way that the operator could perform most of the task while maintaining the connection between the slave and master. However, operators are motivated to decouple and recouple while performing the task, so that they can maintain a comfortable hand/wrist posture. By toggling this connection, operators are able to "scroll through" their workspace, akin to picking up and placing a computer mouse.

Additional information regarding the simulation environment, along with video and download links can be found in Appendix A.

2.5.1 HMD Implementation

The virtual camera linked to the HMD is positioned in the environment such that when seated in a natural position, its location and orientation approximately matches that of the static virtual camera (used for the condition with the desktop monitor).

When viewing the environment through the HMD, a virtual model of the *Touch* haptic device is shown (as seen in Figure 2). This virtual object consists of an imported 3D model⁷, whose virtual position and current end-effector pose is linked to that of the real device. In this way, the operator




Fig. 2. An overview of the simulated environment, where the top figure shows the point of view during operation, and the bottom figure shows this same environment viewed from above. In the top figure, the red border is part of the visual cues (explained in Section 2.7.3). The white grid in the bottom figure has a 1 m x 1 m spacing and is shown here for scale reference. The white lines originating from point 6 show a camera view for the monitor condition (resulting in the view in the top figure). The important objects are numbered for easy reference. 1: The controlled industrial robot. 2: The scoop that is attached to this robot. 3: The zinc bath containing the dross particles. 4: The bins where the dross is deposited into. 5: The steel strip running through the bath. 6: The virtual representation of the *Touch* haptic device. This object was not visible when viewing the scene on the desktop monitor. 7: A window showing how much time has passed. When using the desktop monitor, this same window is shown in the top left corner of the screen.

can intuitively see, and interact with the haptic device when using an HMD, even when it can't be seen directly.

2.5.2 Haptic Feedback

The haptic feedback of the virtual environment is calculated within Unity using a virtual spring. A virtual point is linked to the robot's wrist (point A in Figure 3) using a fixed joint in Unity. When forces are exerted on the robot, this joint "stretches". The OpenHaptics plugin models the stretching of this joint as a linear spring in order to calculate the haptic feedback forces. Additionally, some damping is added to the resulting forces, both on the slave side, as well as the master side. On the slave side, this damping is applied using a linear and angular drag force, applied to all rigid bodies that are part of the controlled robot. On the master side, this damping is calculated using the velocity of the stylus (the pen-like device held by the participant), with a damping coefficient of 0.333 N s/m. The other parameters used in the calculation of these forces are tweaked in such a way that the environment interactions can be felt clearly, without control feeling "sluggish", and without causing instabilities.

3. <https://github.com/ros-industrial/abb/>

4. <https://github.com/siemens/ros-sharp/>

5. <http://obi.virtualmethodstudio.com/>

6. <https://assetstore.unity.com/packages/tools/integration/3d-systems-openhaptics-unity-plugin-134024>

7. <https://grabcad.com/library/omni-phantom-1>



Fig. 3. The scoop attached to the industrial robot. The two red spheres show the locations for which the VFs are activated where A is the wrist joint and B is the scoop. The green volume, consisting of small spheres near location B represent the voxels that are used to approximate the submerged volume of the scoop.

2.6 Design of Haptic Assistance

The HA is supplied in the form of VFs. Guidance forces are applied only in key areas to aid in difficult parts of the task but do not influence movement elsewhere. All forces are exerted on the master side, directly using the haptic device (i.e. not by applying a virtual force to the scoop). The damping is applied to the slave side. This is done by applying a damping force to the scoop based on its linear velocity in the direction of the applied guidance force. The damping constant was tweaked in such a way that no instabilities occurred, without the control feeling "sluggish".

The GVFs above the dross bins apply a force so that the scoop (point B in Figure 3) is pulled towards a point located above the center of the dross deposit bin (e.g. [21]). This force helps in keeping the scoop centered above the bin when depositing the dross. The guidance force is calculated by:

$$F_{VF} = \begin{cases} k(d - d_0)(1 - \frac{d-d_t}{d_g}), & \text{if } d_t < d \leq d_t + d_g \\ k(d - d_0), & \text{if } d \leq d_t \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Here, d is the distance between the tip of the scoop and the guidance point above the dross deposit bin, k is the spring stiffness, d_0 is the rest length of the spring, d_t is the trigger distance from which the force is triggered, and d_g is the gradient distance. This gradient distance allows the force to gradually build-up to the force applied at d_t . This is done so that the operator does not experience a sudden shock because the guidance force jumps to its maximum value.

FRVFs act as an artificial force field, pushing the operator away from dangerous areas (e.g. [40]). The magnitude of this force is calculated using the same equation as the GVFs (Equation 1). However, in this case, d is the distance between the wrist or the scoop (point A, B in Figure 3 respectively), and the boundary of the forbidden region. If both the wrist and the tip of the scoop are within the trigger distance d_t , their resulting forces are summed. The values of these parameters for both the GVF and FRVFs are shown in Table 1.

The GVF in the bath acts as a sort of buoyancy, in which a force is applied upwards, pushing the scoop out of the bath. The magnitude of this force is scaled with the volume of the scoop that is submerged in the bath. This volume is approximated with the use of voxels (volumetric pixels).

TABLE 1
The parameters used for the guidance forces of the VFs.

	k (N/m)	d_0 (m)	d_t (m)	d_g (m)
Bin GVF	2.75	0	0.2	0.1
FRVFs	4.5	0.3	0.3	0




Fig. 4. The visual cues representing the feedback forces and their magnitudes. Here, a) shows the FRVF force cue, b) shows the GVF force cue above the dross deposit bins, c) shows the collision force cue.

Each game update, an array of (1234) evenly spaced voxels (represented by the green balls in Figure 3) is checked to find the fraction of voxels that is below the bath surface. The magnitude of the guidance force is equal to this submerged fraction multiplied by the constant $\alpha = 1.5$. The rationale behind the design of this GVF is that it is important the bath is disturbed as little as possible. Similar to the guidance force, the disturbance also scales with the submerged volume. In this way, the operator receives (indirect) feedback about the bath disturbance and is motivated to keep this disturbance to a minimum. One example of how this is achieved is by using only the front edge of the scoop when gathering (scraping) dross particles in a certain location, before scooping them.

2.7 Design of Visual Feedback

A scene view with the visual cues enabled is seen in Figure 2. In the experimental conditions in which the visual cues are disabled, the visual cues described in this section would not be present. Most of the cues are designed to inform the operator about the haptic feedback currently being experienced. More specifically, they inform the operator of where the force originates from, and visualizes the magnitude of that force. Additionally, they supply the operator with spatial cues, designed to improve spatial, and situational awareness.

One additional visual cue is implemented, which serves purely as a spatial cue. This cue shows the position of the tip of the scoop, projected in the direction of gravity, as seen in Figure 2. A semitransparent, green, dashed line is drawn up until the first surface it hits. At the end of this line, a green circle with a radius of 0.025 m is shown.

The remainder of visual cues are related to the different types of forces that can be experienced. Each of these are visually represented by specialized visual cues and will be explained in the remainder of this section.

2.7.1 Collision Force Visualization

If the scoop collides with another object, a red ring will appear at the place of the collision, as shown in Figure

4c. The radius of this ring is linearly dependent on the magnitude of the collision force, with a minimum of 0.15 m (see the grid in Figure 2 for a scale reference). This scaling is such that at the maximum collision force (Section 2.4) the radius is 1.1 m. Besides informing operators about if and where they are colliding, this cue potentially helps operators make a better judgement about the amount of force they are applying to the environment when colliding [41].

2.7.2 Virtual Fixture Visualization

All VFs are visualized in the environment, as shown in Figure 2. The semitransparent yellow boxes show the boundaries of the FRVFs, while the semitransparent purple areas show the GVFs. This part of the cue is inspired by the work of Giulietti et al. [42], in which the forbidden regions are visualized to improve safety for piloting a remotely piloted vehicle.

Furthermore, if a VF is activated (guidance force is nonzero), an additional cue appears on the location at which this VF is currently active. For the FRVFs and the GVF in the bath, this cue consists of a bright circle, overlaid onto the VF visualization, as shown in Figure 4a. The radius of the circle in the center is dependent on the current magnitude of the guidance force. The outer ring represents the maximum guidance force possible for this VF. If this maximum guidance force is reached and the whole circle is filled, it will start to blink. In this way, operators are informed of where the force comes from and what its magnitude is. Additionally, operators are (indirectly) informed of how close they are to the boundary of the virtual fixture. This visual cue is in part inspired by the work of Nakazawa et al. [27], in which proximity to the to a cone-shaped boundary in a deep and narrow surgery task is visualized using a semicircle on the edges of the screen. Furthermore, visualizing the operational boundaries, and the controlled systems' current relation to them, has been shown to bring benefits in automotive [29], and aviation domain [20]. This cue is designed to bring a similar visual cue, but applied to the teleoperation of an industrial robot.

The GVF above the bins is different, as it consists of a single point. The GVF itself is visualized by a purple sphere, centered onto the guidance point that it represents. The radius of this sphere is the trigger distance described in Section 2.6. Unlike the other two VF cues, this cue doesn't show the magnitude of the applied force, but simply lights up an area around the point on the scoop on which the GVF acts, as shown in Figure 4b. This area lights up slightly before the scoop gets close enough to activate it, so that the guidance force and its direction can be anticipated.

The appearance of these visual cues and the way in which they is coded is further explained in Appendix A and Appendix B.

2.7.3 Workspace Visualization

During teleoperation, there are two workspaces to take into account: that of the industrial robot, and that of the Touch haptic device. Reaching the limits of either of the two workspaces results in a force experienced by the operator.

Because of this, both of these forces are visualized using visual cues.

If the workspace limits of the robot are reached, the force feedback behaves the same as when a collision with the environment occurs. The controlled point can not pull the robot any further, and so a force pulls it back within the workspace limits. The workspace for the robot is visualized to inform the operator of this limit. The visualization of this resulting force behaves in the same way as described for the FRVFs and GVF in the bath discussed in the previous section, and the design choices here follow the same rationale.

The limits of the haptic device feel different, as the stylus can't move any further. Nonetheless, it can be difficult to recognize this difference during operation. To draw attention to the fact that the devices' workspace limits have been reached, a red border shows up in the direction in which the limits have been reached, or are about to be reached. Here the up, down, left and right limits correspond to the lighting up of the top, bottom, left and right border of the screen, respectively. When the back and front limits (close to and away from the operator) are reached, all four borders of the screen light up, as seen in Figure 2. When the HMD was used, this cue is shown as a floating square in front of the operator, as if wearing some sort of spectacle. The size of this square is such that it doesn't block the central vision while remaining visible when looking at the center of the display.

This visual cue is inspired by the work of Galambos et al. [43], in which the peripheral field of vision is used to display a visualization of the grasping force. The study shows that humans can extract information using peripheral vision independently from the central vision, as long as the information is easy to interpret.

2.8 Experimental Design

The effect of the two independent variables was evaluated in a counterbalanced 2 (Display: Monitor vs. HMD) \times 2 (Cues: With vs. Without) within-subjects design. In other words, there is a total of four experimental conditions, which for ease of reference are abbreviated as:

- *MN*: Monitor display, and No visual cues
- *MC*: Monitor display, and with visual Cues
- *HN*: HMD, and No visual cues
- *HC*: HMD, and with visual Cues

These conditions were ordered according to an incomplete counterbalanced measures design. Prior to the experiment participants read and signed an informed consent form (Appendix D). They received a separate document detailing this consent form (Appendix C). This document gives an overview of the experiment, its purpose, risks and its procedure. Afterwards, they received a training session of about 10 minutes, in which they were able to toggle the presence of each of the haptic and visual cues, and switch between displays. This control is enabled through a simple user interface (UI) with buttons. Additionally, this UI provided information boxes about the purpose

and functionality of each cue. As they practiced with the interface, they are informed to also try out the boundary conditions defined in Section 2.4. Furthermore, participants are instructed to minimize the collision forces and the disturbance to the bath. Once the participant was confident in their ability to use the interface, and had a clear understanding of the experiment, a short questionnaire is filled out, in which the participant's demographics were collected.

After this, participants started with the first condition, where the current experimental condition was displayed on the screen. The current repetition only started when the haptic device is first connected to the robot (as explained in Section 2.5).

First, a practice repetition of five minutes was performed, in which participants familiarized themselves with the current interface, and practiced their scooping strategy. Participants were requested to take this repetition as seriously as a real repetition, so that they are able to familiarize with all aspects of the operation. After this, the participant performed two real repetitions, which together will be used for the calculations of the dependent measures. If one of the repetitions had failed (due to violating one of the BCs), only that repetition was stopped. After each repetition, participants were informed of the amount of dross they had removed, after which a short break (of at least one minute) was started. Participants were allowed to take longer breaks if they wanted to, and were asked only to continue when they felt rested. After all three repetitions, the participants were asked to fill out a *Van der Laan questionnaire* [44] to assess the usefulness and satisfaction of the interface as a whole. These steps were repeated for each of the experimental conditions (four times in total).

After the last Van der Laan questionnaire was completed, participants were asked to fill out one *final questionnaire*. This questionnaire consisted of four questions. First, they were asked what was difficult about this task (Q1). The next question asked what experimental conditions they liked the most (Q2), and why (Q3). Lastly, they were asked if they had any additional remarks (Q4). The total experiment, including the training and filling out the questionnaires took approximately 2 hours per participant. All questionnaires are displayed in Appendix E.

2.9 Dependent Measures

All data was recorded within Unity game engine and was logged every frame update. The two repetitions performed for each condition are combined for the calculation of the dependent measures. If a repetition is shortened because the BCs were violated, the data for this repetition is still included, but that trial simply lasted shorter. In other words, failing a repetition shortened the amount of time that participants had to remove dross from the bath. The calculated dependent measures have been categorized into performance, safety, and user acceptance.

2.9.1 Performance

- *Percentage of Dross Removed (%)*. Dross is considered removed, only when the particles have been de-

posited into one of the two bins. The total amount of dross removed over the two repetitions is divided by the total amount of dross present in those repetitions (1500 particles). This is used as the main performance measure.

- *Average Bath Disturbance (-)*. The bath disturbance is approximated by multiplying the fraction of the submerged volume (as used by the GVF in the bath, described in Section 2.6) with the velocity of the end of the scoop (point B, shown in Figure 3). The resulting value is averaged over time. Participants were instructed that for this task, it is important to minimize this measure by submerging the scoop only as much as is necessary, especially while making large motions through the bath. The resulting scores are used as a measure for the performance, as a high disturbance will lead to unsatisfactory results in the galvanizing process, regardless of the amount of dross that is removed. Additionally, this measure provides an indication of the accuracy with which the task can be performed, as minimizing this measure requires operators to maintain an accurate distance to the liquid level.

2.9.2 Safety

- *Number of repetitions failed*. Repetitions fail and are stopped if the boundary conditions are violated. Besides counting this number, the failure of a repetition is reflected in the performance measures as well, as this failure leads to participants having less time to perform the operation.
- *Peak collision Force (N)*. Participants are instructed to minimize the collision forces. The peak collision force gives an indication of how close they got to failing the trial as a result of environment collisions. The peak collision force is the maximum exerted force over both repetitions. This value is used as the main safety measure.
- *Minimum distance to steel strip (m)*. Collision with the steel strip results in big financial losses. For this reason, if this occurred during the experiment, that repetition fails and is stopped immediately. As such, getting close to the steel strip is risky, and the minimum distance to the steel strip gives an indication of the amount of risk that of the operation.

2.9.3 User Acceptance

- *Preferred Condition*. In the final questionnaire, participants are asked which experimental condition they preferred.
- *Van der Laan Questionnaire*. The user acceptance of the interface is assessed using the Van der Laan acceptance questionnaire [44], by evaluating participant reported usefulness and satisfaction scores. Participants were specifically instructed to rate the interface as a whole (the cues, the display, the haptic interface, and the haptic feedback and guidance).

2.10 Statistical Analyses

The statistical analyses were done using a two-way repeated measures analysis of variance (RM ANOVA) [45]. When

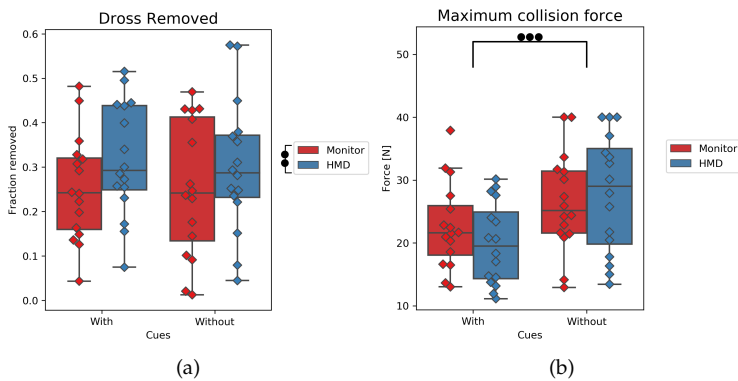


Fig. 5. (a) Percentage of dross removed (main performance measure), (b) Maximum collision force (main safety measure). The marks denote significance, where $\bullet p \leq 0.05$, $\bullet\bullet p \leq 0.01$, $\bullet\bullet\bullet p \leq 0.001$

interaction effects are significant, a post-hoc analysis is done to re-evaluate the main effects. For the parametric data, this is done using an independent samples t-test, and for the nonparametric data, the Wilcoxon signed-rank test is used. Data were checked for normality using the Shapiro-Wilk test [46] on each combination of the independent variables. Data sets that did not pass this test were transformed using the Aligned Rank Transform, as described by Wobbrock et al. [47], using the ARTool⁸.

Results were regarded as statistically significant when $p \leq 0.05$.

3 RESULTS

In this section, the experimental results are presented in accordance with the dependent measures as explained in Section 2.9. Table 2 shows the means and standard deviations for all dependent measures, along with the results of the two way RM ANOVA. The most important results are discussed in the text in this section, though for the full details of the statistical analysis, the reader is referred to Table 2. As only one measure required post-hoc pairwise analysis (satisfaction score), its results are presented in the text in Section 3.3. The results of the Shapiro Wilk test for all data sets are reported in Appendix F.

3.1 Performance

For the percentage of dross removed, there was a significant main effect of the display, $F(1, 15) = 9.68$, $p = 0.007$ (Table 2; Figure 5a). This effect was such that with the HMD a significantly higher percentage of dross was removed than with the monitor. No statistically significant main effect of the cues ($p = 0.74$), and no statistically significant interaction effect ($p = 0.72$) was found for this measure. Furthermore, no significant effects were found for the average bath disturbance (Table 2). However, it is noteworthy that the means of the bath disturbance show a slight, but insignificant decrease with the implementation of the visual cues ($p = 0.10$).

8. <http://depts.washington.edu/acelab/proj/art/index.html>

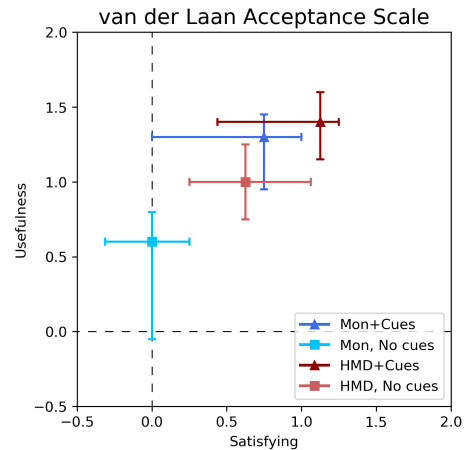


Fig. 6. Van der Laan acceptance scores [44]; on the horizontal axis the satisfying dimension and on the vertical axis the usefulness dimension. Scores are in the range of (-2,2). The center points are the median values, and the colored lines show 25% percentile ranges. Additional details are stated in Table 2.

3.2 Safety

Out of the combined total of repetitions of all participants and all conditions (128 repetitions), a total of 6 have failed due to a violation of one of the boundary conditions. All failed repetitions occurred during the conditions without visual cues, where 3 occurred with the HMD (HN) and 3 occurred with the desktop monitor (MN). Only in one of these repetitions the failure was caused by a collision with the steel strip, which occurred during the MN condition.

For the peak collision force, a significant effect of the cues is found $F(1, 15) = 22.19$, $p = .0003$ (Table 2; Figure 5b). This effect is such that the peak collision force was significantly lower with visual cues compared to without them. However, no such effect was found for the type of display ($p = 0.61$), and no significant interaction effect was found ($p = 0.29$).

Furthermore, no significant effects were found for the minimum distance to the steel strip (Table 2).

3.3 User Acceptance

Out of the 16 participants, 12 participants preferred the HC condition (Section 2.8). The other 4 participants were evenly divided over the MC and HN conditions. Furthermore, both the implementation of the HMD and the visual cues significantly improve both the satisfaction and the usefulness scores (Table 2; Figure 6).

However, the interaction effect of the satisfaction score was also significant, $F(1, 15) = 5.03$, $p = 0.04$. A pairwise post-hoc comparison reveals a significant effect of the visual cues, only when the HMD was not used (MC-MN: $p = 0.0048$, HC-HN: $p = 0.071$). Likewise, the effect of the implementation of the HMD is only significant when the visual cues are not implemented (MN-HN: $p = 0.0015$, MC-HC: $p = 0.173$).

TABLE 2

Means (M), standard deviations (SD), and results of the two way repeated measures ANOVA for each dependent measure. Significant effects ($p \leq 0.05$) are printed in **bold**. Condition abbreviations are described in Section 2.8.

Measures	Conditions				RM ANOVA, F(1,15)			
	MN	MC	HN	HC	Cues	Display	Interaction	
Performance								
Percentage Dross Removed (%)	M	25.28	25.36	30.12	31.73	$p = 0.7425$	$p = \mathbf{0.0072}$	$p = 0.7174$
	SD	14.89	11.54	14.53	12.46	$F = 0.11$	$F = 9.68$	$F = 0.13$
Average Bath Disturbance (volume fraction $\cdot m/s$)	M	$2.21 \cdot 10^{-2}$	$1.72 \cdot 10^{-2}$	$1.96 \cdot 10^{-2}$	$1.73 \cdot 10^{-2}$	$p = 0.1043$	$p = 0.1882$	$p = 0.3589$
	SD	$1.68 \cdot 10^{-2}$	$0.96 \cdot 10^{-2}$	$1.39 \cdot 10^{-2}$	$1.03 \cdot 10^{-2}$	$F = 2.99$	$F = 1.90$	$F = 0.95$
Safety								
Number of failed trials (Count)		3	0	3	0			
Peak collision force (N)	M	26.40	22.62	27.89	19.89	$p = \mathbf{0.0003}$	$p = 0.6116$	$p = 0.2917$
	SD	7.59	6.61	9.13	6.30	$F = 22.19$	$F = 0.27$	$F = 1.19$
Min. distance to steel strip (m)	M	0.57	0.71	0.64	0.63	$p = 0.1399$	$p = 0.3518$	$p = 0.5342$
	SD	0.20	0.27	0.30	0.23	$F = 2.43$	$F = 0.92$	$F = 0.41$
User Acceptance								
Preferred Condition (Count)		0	2	2	12			
Satisfaction score* (-2,2)	M	0.28	1.13	0.98	1.25	$p = \mathbf{0.0012}$	$p = \mathbf{0.0010}$	$p = \mathbf{0.0405}$
	SD	0.52	0.72	0.49	0.69	$F = 15.84$	$F = 16.76$	$F = 5.03$
Usefulness score* (-2,2)	M	-0.09	0.50	0.70	0.81	$p = \mathbf{0.0100}$	$p < \mathbf{0.0001}$	$p = 0.1316$
	SD	0.76	0.56	0.48	0.58	$F = 8.68$	$F = 42.61$	$F = 2.54$

* Shapiro-Wilk test violated, data transformed using aligned rank transformation [47].

4 DISCUSSION

This study proposes two methods to provide the operators with additional information through the visual channel with the goal of complementing HA enabled teleoperation. Both proposed methods (implementation of visual cues, and the use of an HMD instead of a desktop monitor) provide spatial cues. Additionally, the visual cues provide information regarding the HA, and other forces felt during operation. In contrast, the HMD provides operators with a better sense of immersion. To understand the efficacy of these methods in the proposed interface their implementation is evaluated based on performance, safety and user acceptance.

All hypotheses were found to be partially confirmed and partially rejected. H1 was partially confirmed in that user acceptance was improved by the implementation of visual cues, though safety was only partially improved, and the hypothesis of the improvement in performance was rejected. The hypothesis H2 was confirmed in that the use of an HMD improved user acceptance, although the hypothesized improvement in performance was only partially confirmed.

The hypothesis H3 of the absence of an interaction effect was confirmed for all measures, except for the satisfaction score. This suggests that, as hypothesized, the two visual feedback methods have some unique benefits, and (when an effect is found) result in improvements regardless of the presence of the other method. The exception to this result is discussed below, in Section 4.4.

The implications of these results and their relation to previous research will be discussed in the remainder of this

section. The answers to the final questionnaire have been used to motivate the discussion of these results. This has led to the supplementary analysis and discovery of some unforeseen effects, which are discussed in Section 4.1.

4.1 Supplementary Analysis: The Effect of the Display Type on Manipulability

In the final questionnaire, the majority (9 participants) reported the control of the robot, specifically controlling the orientation of the scoop, to be a major source of difficulty (Q1). A likely reason for this is the asymmetry in the workspace and mechanical structure between the haptic device and the robot. This provokes operators to frequently toggle the master/slave connection, as a result of which the orientation of the scoop and that of the pen rarely coincide. Participants were also asked to explain their choice of preferred experimental condition (Q3). The two most common reasons given for the preference of the HMD are that (1) operators experienced a better spatial/situational awareness, making it easier to navigate around difficult areas (stated by 9 participants), and (2) operators reported an increased sense of immersion, felt more present, and experienced the control of the robot to feel more intuitive/natural (stated by 5 participants). These comments suggest that potentially, these problems in manipulability were mitigated when using an HMD. To verify these statements and gain additional insights on the results discussed in this section, an extra data analysis was done on (1) the average velocity, and (2) the average amount of dross particles per dross deposit.

4.1.1 Average Velocity

The two-way RM ANOVA of the average velocity reveals that indeed there was a statistically significant main effect

of the display, $F(1, 15) = 5.71$, $p = 0.0304$, such that the average velocity is higher when using an HMD, compared to using a desktop monitor. No significance was found for the main effect of the cues ($F(1, 15) = 2.81$, $p = 0.1145$), or the interaction effect ($F(1, 15) = 0.08$, $p = 0.7856$). The means \pm the standard deviations for the MN, MC, HN, HC conditions were 0.45 ± 0.08 , 0.44 ± 0.09 , 0.47 ± 0.09 , and 0.45 ± 0.08 m/s respectively.

This result suggests that with the HMD, operators had an easier time navigating the environment, and were able to move slightly (but significantly) faster as a result.

4.1.2 Average Scoop Size

Appropriately orienting the scoop is absolutely critical for scooping dross from the bath. Because of this, the reported difficulty in controlling the orientation of the scoop will likely result in a lower number of scoops and/or less dross particles being collected each scoop.

These results are calculated by finding the timestamps at which changes took place in the amount of dross in the bins, and in the bath. Using this data, the task can be divided into the subtasks of dross scooping and dross dumping, yielding the total number of scoops for each condition. Dividing the total amount of dross particles collected by the number of scoops gives the average amount of dross particles per scoop.

In regards to the total number of scoops, no significant main effect for cues ($F(1, 15) = 0.48$, $p = 0.5000$) or the display type ($F(1, 15) = 3.98$, $p = 0.0644$) were found. Furthermore, no interaction effects are found ($F(1, 15) = 0.00$, $p = 1.00$). The means \pm the standard deviations for the MN, MC, HN, HC conditions were 24.50 ± 11.28 , 25.88 ± 7.30 , 26.88 ± 10.02 , and 28.25 ± 6.75 particles respectively.

Conversely, a statistically significant main effect of the display was found regarding the average amount of particles collected per scoop, $F(1, 15) = 21.11$, $p = 0.0004$, such that the average scoop size is larger when using an HMD, compared to using a desktop monitor. No significance was found for the main effect of the cues ($F(1, 15) = 0.15$, $p = 0.7081$), or the interaction effect ($F(1, 15) = 0.02$, $p = 0.8918$). The means for the MN, MC, HN, HC conditions were 14.34, 14.75, 16.75, and 16.95 particles respectively.

These results suggest that the HMD helps operators in controlling the scoop in such a way that more particles can be scooped per motion. Relating these results to the questionnaire answers discussed above, it is a likely assumption that the use of an HMD improves the manipulability of the proposed interface.

One possible reason for this could be that because participants are more immersed, they forget about the haptic interface and perceive the connection more as "grabbing" the scoop, rather than linking the two devices. This makes the interaction more intuitive and could explain these improvements. Another reason could be that because of the improvements in depth perception participants could make

a better judgment about the orientation and position of the scoop.

4.2 Effects of Visual cues on Safety

The visual cues were found to significantly improve the safety of the operation, such that the peak collision force was significantly lower, partially confirming the first hypothesis. This indicates that visual cues are important in improving the safety of the operation. This statement is supported by the fact that zero of the six failed repetitions occurred when the visual cues were present. Similar improvements in safety have been found in previous research in the aviation [20] [31], automotive [29], and (narrow) surgery [27] domains. A likely reason as to why these improvements are found is that participants were more aware of what forces are present, and had a better understanding of where these forces come from (which was a common statement in the comments of the final questionnaire).

That said, this result is not reflected in the minimum distance to the steel strip, which opposes the results in stated previous research. A logical reason for this is that in all tasks in this previous research, getting close to the forbidden regions provides no direct benefits and mostly induces risk. In contrast, in this task getting close to the forbidden regions (even penetrating them on some occasions) allows operators to collect more of the dross particles, increasing their performance scores. The final questionnaire revealed that operators felt like they were more aware of the perceived dangers with the presence of visual cues. This is likely the reason participants did not keep a greater distance to the steel strip and were willing to accept the risks that come with that.

4.3 Effects on Performance

In contrast to the first hypothesis, task performance was not significantly affected by the addition of visual cues. Borst et al. [20] evaluated the implementation of visual cues in addition to HA, finding evidence it helped operators to more accurately follow an optimal trajectory. It was expected that similarly, the addition of the visual cues would result in a reduction of the average bath disturbance. Although a slight decrease was found, the effect was not significant, suggesting that the visual cues did not help the participants in recognizing how much disturbance they are causing to the bath. Additionally, we found relatively large standard deviations for these datasets (Table 2). This variability could be explained by the fact that participants did not receive any direct feedback about how much disturbance they had caused for each trial (as opposed to the percentage of dross removed). Most likely, operators were unaware of the significance of this aspect of the experiment, despite it being explicitly stated during the training session.

The absence of an effect for the percentage of dross removed could be caused by participants having a better sense of danger, and as a result being more careful in their navigation around the environment (Section 4.2). This could prevent the predicted increase of performance, even though

operators are supplied with additional spatial cues.

The second hypothesis, predicting an improvement in performance with the use of an HMD, was only partially confirmed. A significant effect was found for the main performance measure (percentage of dross removed). This indicates that the use of an HMD helps operators in efficiently controlling the robot such that more dross can be removed. This result confirms previous research, in which performance enhancements were found when depth perception is an important factor [36], and when the interface is not intuitive [37]. The improvements in performance are more clearly expressed by the improved manipulability (as described in Section 4.1).

Even still, although hypothesized, such an effect was not found in the results of the average bath disturbance. This indicates that even though the HMD helps operators in positioning and orienting the scoop, it did not help in accurately keeping the scoop at just the right height, to minimize bath disturbance. Additionally, as explained above, the large variability in these results could be caused by a lack of clear feedback regarding the amount of disturbance caused.

4.4 Effects on user Acceptance

As hypothesized, both the visual cues (H1), as well as the HMD (H2) result in a significantly improved user acceptance, for both the satisfaction, as well as the usefulness score. This is further supported by the scores of the preferred condition, as seen in Table 2. This result is in accordance with previous studies, in which similar improvements in user acceptance were found as a result of the implementation of visual cues [31] [20] [27], and the use of an HMD [37]. This study shows that the results found in these studies can be extended to their use in the HA teleoperation of an industrial robot.

However, in contrast to the hypothesis (H3), a significant interaction effect was found for the satisfaction score. The pair-wise comparison revealed that when one of the methods was already implemented, participants did not perceive the interface as more satisfying with the addition of the other method. A possible explanation for this is that the visual cues can be distracting, or obstruct the line of sight, as two participants have commented was sometimes the case. This could indicate that when the information provided by these cues is not perceived as useful, user acceptance might not increase further, as the cues might only serve as a distraction. This is in analogy with the work of Ho et al. [31], in which it was found that in some situations, the visual cues would cause clutter resulting in a decrease in user acceptance and even decreased performance.

4.5 Future Work

In this study, no evaluations have been done concerning the operator workload. On the one hand, overloading operators with excessive information increases workload and could diminish the benefits that have been found [48]. On the other hand, visual cues in HA have been shown to reduce workload, as they can make the haptic information easier to interpret [49] [27]. As such, investigating how this relates

to the two methods suggested in this research is far from trivial and could be an interesting topic for future research. Furthermore, since the visual cues have been implemented simultaneously, this study provides little insight into the effects of the individual cues. Obtaining this insight could lead to a more effective, specialized set of visual cues.

One of the more unexpected results is the improved manipulability found when using an HMD. Participants stated feeling more immersed, reporting a sense of embodiment making the control through the haptic interface feel more natural. This synergy could also prove to be beneficial in other tasks requiring dexterous manipulation, and is worth exploring further.

Additionally, it would also be interesting to investigate visual feedback for implementations of HA other than VFs, such as Haptic Shared Control ([15]). A full adaptation would require modifying the VF specific visual cues, which will depend heavily on the type of HA used and application it is used for. Nonetheless, the other visual cues and the use of an HMD can quite easily be implemented in any other type of HA teleoperation interface for industrial robots, and evaluating their effects would be an interesting extension of this research.

Finally, in the context of the proposed use case of dross removal, we suggest the implementation of both visual feedback methods. However, some additional research should be done on the workload and how these results translate to a real robot. The implementation of both methods will require multiple sensors and virtualization ([50]) techniques. This could prove to be challenging as additional problems such as sensor noise and delays could diminish the benefits observed in this study. Before implementation in a real dross removal environment, these challenges should be investigated.

5 CONCLUSION

In this study, the implementation of a set of visual cues, as well as the use of an HMD have been found to provide benefits for the purpose of the HA teleoperation of an industrial robot using VFs. Both methods, and the interaction between them, were evaluated in a two-way human factors experiment for the task of dross removal in a simulated environment. From the results of this experiment, the following is concluded:

- The visual cues improve the safety of the operation, helping to prevent collisions with high forces, and with dangerous obstacles. These effects were found regardless of the implementation of an HMD.
- The use of an HMD compared to the use of a desktop monitor increases the main task performance, without compromising the safety and accuracy of the operation.
- Supplementary analysis shows that the manipulability of the system was increased with the implementation of the HMD as participants were able to navigate more quickly and manipulate the system more effectively.

- The performance and manipulability improvements caused by the use of an HMD are present regardless of the implementation of the visual cues.
- In general, system acceptance improved with the implementation of both of the proposed methods. However, when one of the methods is already present, the implementation of the other method does not further increase the satisfaction scores.

These results indicate that although both methods provide benefits, they do so in different aspects. Task performance mostly benefits from the use of an HMD, whereas safety mostly benefits from the use of visual cues. Moreover, the results for user acceptance indicate that although both methods improve user acceptance, combining them might not cause further improvements. As such, interface designers should be careful to consider the necessity of including each of the proposed visual feedback methods.

This study takes an important step towards gaining a better understanding of the importance of visual feedback design in HA teleoperation. Furthermore, this study adds to the limited amount of work that has investigated this interaction and provides new evidence to support the idea that this synergy is worth exploring further.

REFERENCES

- [1] O. Khatib, "The new robotics age: Meeting the physical interactivity challenge," in *Symposium on Robot Design, Dynamics and Control*. Springer, 2016, pp. 17–18.
- [2] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human factors*, vol. 39, no. 2, pp. 230–253, 1997.
- [3] J. D. Lee, "Review of a pivotal human factors article: 'humans and automation: Use, misuse, disuse, abuse'," *Human Factors*, vol. 50, no. 3, pp. 404–410, 2008, pMID: 18689046. [Online]. Available: <https://doi.org/10.1518/001872008X288547>
- [4] J. B. van Erp, M. Duistermaat, C. Jansen, E. Groen, and M. Hoedemaeker, "Tele-presence: Bringing the operator back in the loop," HUMAN FACTORS RESEARCH INST TNO SOESTERBERG (NETHERLANDS), Tech. Rep., 2006.
- [5] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared control in haptic systems for performance enhancement and training," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 1, pp. 75–85, 2006.
- [6] J.-y. Lee and S. Payandeh, *Haptic Teleoperation Systems*. Springer, 2015.
- [7] J. G. Wildenbeest, D. A. Abbink, C. J. Heemskerk, F. C. Van Der Helm, and H. Boessenkool, "The impact of haptic feedback quality on the performance of teleoperated assembly tasks," *IEEE Transactions on Haptics*, vol. 6, no. 2, pp. 242–252, 2012.
- [8] M. J. Massimino and T. B. Sheridan, "Teleoperator performance with varying force and visual feedback," *Human Factors*, vol. 36, no. 1, pp. 145–157, 1994, pMID: 8026837. [Online]. Available: <https://doi.org/10.1177/001872089403600109>
- [9] J. Y. C. Chen, E. C. Haas, and M. J. Barnes, "Human performance issues and user interface design for teleoperated robots," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 6, pp. 1231–1245, 2007.
- [10] M. A. Goodrich, J. W. Crandall, and E. Barakova, "Teleoperation and beyond for assistive humanoid robots," *Reviews of Human factors and ergonomics*, vol. 9, no. 1, pp. 175–226, 2013.
- [11] J. Draper, W. Moore, J. Herndon, and B. Weil, "Effects of force reflection on servomanipulator task performance," Oak Ridge National Lab., TN (USA), Tech. Rep., 1986.
- [12] B. Hannaford, L. Wood, D. A. McAfee, and H. Zak, "Performance evaluation of a six-axis generalized force-reflecting teleoperator," *IEEE transactions on systems, man, and cybernetics*, vol. 21, no. 3, pp. 620–633, 1991.
- [13] H. S. Vitense, J. A. Jacko, and V. K. Emery, "Multimodal feedback: an assessment of performance and mental workload," *Ergonomics*, vol. 46, no. 1-3, pp. 68–87, 2003.
- [14] H. Boessenkool, D. A. Abbink, C. J. Heemskerk, and F. C. van der Helm, "Haptic shared control improves tele-operated task performance towards performance in direct control," in *2011 IEEE World Haptics Conference*. IEEE, 2011, pp. 433–438.
- [15] D. A. Abbink, M. Mulder, and E. R. Boer, "Haptic shared control: smoothly shifting control authority?" *Cognition, Technology & Work*, vol. 14, no. 1, pp. 19–28, 2012.
- [16] D. A. Abbink, T. Carlson, M. Mulder, J. C. de Winter, F. Aminravan, T. L. Gibo, and E. R. Boer, "A topology of shared control systems—finding common ground in diversity," *IEEE Transactions on Human-Machine Systems*, no. 99, pp. 1–17, 2018.
- [17] M. Mulder, D. A. Abbink, and E. R. Boer, "Sharing control with haptics: Seamless driver support from manual to automatic control," *Human Factors*, vol. 54, no. 5, pp. 786–798, 2012, pMID: 23156623. [Online]. Available: <https://doi.org/10.1177/0018720812443984>
- [18] L. B. Rosenberg, "Virtual fixtures: Perceptual tools for telerobotic manipulation," in *Proceedings of IEEE Virtual Reality Annual International Symposium*, 9 1993, pp. 76–82.
- [19] J. J. Abbott, P. Marayong, and A. M. Okamura, "Haptic virtual fixtures for robot-assisted manipulation," in *Robotics Research*, S. Thrun, R. Brooks, and H. Durrant-Whyte, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 49–64.
- [20] C. Borst, F. Grootendorst, D. Brouwer, C. Bedoya, M. Mulder, and M. Van Paassen, "Design and evaluation of a safety augmentation system for aircraft," *Journal of Aircraft*, vol. 51, no. 1, pp. 12–22, 2013.
- [21] S. Payandeh and Z. Stanicic, "On application of virtual fixtures as an aid for telemanipulation and training," in *Proceedings 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2002*. IEEE, 2002, pp. 18–23.
- [22] J. van Oosterhout, C. J. Heemskerk, H. Boessenkool, M. R. De Baar, F. C. van der Helm, and D. A. Abbink, "Haptic assistance improves tele-manipulation with two asymmetric slaves," *IEEE transactions on haptics*, vol. 12, no. 2, pp. 141–153, 2018.
- [23] H. Boessenkool, D. A. Abbink, C. J. M. Heemskerk, F. C. T. van der Helm, and J. G. W. Wildenbeest, "A task-specific analysis of the benefit of haptic shared control during telemanipulation," *IEEE Transactions on Haptics*, vol. 6, no. 1, pp. 2–12, 1 2013.
- [24] M. Itoh, F. Flemisch, and D. Abbink, "A hierarchical framework to analyze shared control conflicts between human and machine," *IFAC-PapersOnLine*, vol. 49, no. 19, pp. 96 – 101, 2016, 13th IFAC Symposium on Analysis, Design, and Evaluation of Human-Machine Systems HMS 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2405896316320584>
- [25] F. Flemisch, J. Kelsch, C. Löper, A. Schieben, and J. Schindler, "Automation spectrum, inner/outer compatibility and other potentially useful human factors concepts for assistance and automation," *Human Factors for assistance and automation*, no. 2008, pp. 1–16, 2008.
- [26] T. Lam, M. Mulder, M. Van Paassen, J. Mulder, and F. Van der Helm, "Force-stiffness feedback in uav tele-operation with time delay," in *AIAA Guidance, Navigation, and Control Conference*, 2009, p. 5977.
- [27] A. Nakazawa, K. Nanri, K. Harada, S. Tanaka, H. Nukariya, Y. Kurose, N. Shono, H. Nakatomi, A. Morita, E. Watanabe *et al.*, "Feedback methods for collision avoidance using virtual fixtures for robotic neurosurgery in deep and narrow spaces," in *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, 2016, pp. 247–252.
- [28] K. Christoffersen and D. D. Woods, "How to make automated systems team players," in *Advances in human performance and cognitive engineering research*. Emerald Group Publishing Limited, 2002, pp. 1–12.
- [29] W. Vreugdenhil, "Complementing automotive haptic shared control with visual feedback for obstacle avoidance," 2019.
- [30] J. Lee, J. Hoffman, H. Stoner, B. Seppelt, and M. Brown, "Application of ecological interface design to driver support systems," in *Proceedings of IEA*, vol. 2006. Citeseer, 2006, p. 16th.
- [31] V. Ho, C. Borst, M. M. van Paassen, and M. Mulder, "Increasing acceptance of haptic feedback in uav teleoperation by visualizing force fields," in *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2018, pp. 3027–3032.

- [32] R. J. Kuiper, D. J. Heck, I. A. Kuling, and D. A. Abbink, "Evaluation of haptic and visual cues for repulsive or attractive guidance in nonholonomic steering tasks," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 5, pp. 672–683, 2016.
- [33] E. D. Ragan, R. Kopper, P. Schuchardt, and D. A. Bowman, "Studying the effects of stereo, head tracking, and field of regard on a small-scale spatial judgment task," *IEEE transactions on visualization and computer graphics*, vol. 19, no. 5, pp. 886–896, 2012.
- [34] G. Burdea, P. Richard, and P. Coiffet, "Multimodal virtual reality: Input-output devices, system integration, and human factors," *International Journal of Human-Computer Interaction*, vol. 8, no. 1, pp. 5–24, 1996.
- [35] T. B. Sheridan, "Defining our terms," *Presence: Teleoperators & Virtual Environments*, vol. 1, no. 2, pp. 272–274, 1992.
- [36] O. Liu, D. Rakita, B. Mutlu, and M. Gleicher, "Understanding human-robot interaction in virtual reality," in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2017, pp. 751–757.
- [37] D. Whitney, E. Rosen, E. Phillips, G. Konidaris, and S. Tellex, "Comparing robot grasping teleoperation across desktop and virtual reality with ros reality," in *Robotics Research*. Springer, 2020, pp. 335–350.
- [38] B. Xi, S. Wang, X. Ye, Y. Cai, T. Lu, and R. Wang, "A robotic shared control teleoperation method based on learning from demonstrations," *International Journal of Advanced Robotic Systems*, vol. 16, no. 4, p. 1729881419857428, 2019.
- [39] X. Hao, L. Bai, C. Yue, X. Xu, Z. Yin, and M. Zhang, "The research and application of zinc dross removing robot in galvanized production line," 2018.
- [40] J. J. Abbott, "Virtual fixtures for bilateral telemanipulation," Ph.D. dissertation, Citeseer, 2005.
- [41] T. Horeman, M. D. Blikkendaal, D. Feng, A. van Dijke, F. Jansen, J. Dankelman, and J. J. van den Dobbelsteen, "Visual force feedback improves knot-tying security," *Journal of surgical education*, vol. 71, no. 1, pp. 133–141, 2014.
- [42] F. Giuletta, L. Pollini, and G. Avanzini, "Visual aids for safe operation of remotely piloted vehicles in the controlled air space," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 230, no. 9, pp. 1641–1654, 2016.
- [43] P. Galambos, A. Róka, G. Sörös, and P. Korondi, "Visual feedback techniques for telemanipulation and system status sensualization," in *2010 IEEE 8th International Symposium on Applied Machine Intelligence and Informatics (SAMII)*. IEEE, 2010, pp. 145–151.
- [44] J. D. Van Der Laan, A. Heino, and D. De Waard, "A simple procedure for the assessment of acceptance of advanced transport telematics," *Transportation Research Part C: Emerging Technologies*, vol. 5, no. 1, pp. 1–10, 1997.
- [45] J. Verma, *Repeated measures design for empirical researchers*. John Wiley & Sons, 2015.
- [46] S. S. Shapiro and M. B. Wilk, "An analysis of variance test for normality (complete samples)," *Biometrika*, vol. 52, no. 3/4, pp. 591–611, 1965.
- [47] J. O. Wobbrock, L. Findlater, D. Gergle, and J. J. Higgins, "The aligned rank transform for nonparametric factorial analyses using only anova procedures," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2011, pp. 143–146.
- [48] E. B. Entin, E. E. Entin, and D. Serfaty, "Optimizing aided target-recognition performance," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 40, no. 4, pp. 233–237, 1996. [Online]. Available: <https://doi.org/10.1177/154193129604000419>
- [49] S. D. Stigter, M. Mulder, and M. Van Paassen, "Design and evaluation of a haptic flight director," *Journal of guidance, control, and dynamics*, vol. 30, no. 1, pp. 35–46, 2007.
- [50] B. P. Vagvolgyi, W. Pryor, R. Reedy, W. Niu, A. Deguet, L. L. Whitcomb, S. Leonard, and P. Kazanzides, "Scene modeling and augmented virtuality interface for telerobotic satellite servicing," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4241–4248, 2018.

A

Unity Scene

This appendix contains additional information regarding the simulated environment, referred to as the *Unity scene*. This information (together with the scientific paper) provides guidance for reproducing this experiment and its environment, as well as additional information for the reader that is interested the simulation environment and how it is set-up within the Unity game engine.

A.1. Video Demonstration

To get a more comprehensive view of the experiment and the virtual environment in which it takes place, a video demonstration can be viewed on the following url:

<https://vimeo.com/448282487>

This video shows one full condition of the experiment (three repetitions in total). Additionally, the starting questionnaire and the van der Laan questionnaire are shown and filled in as a demonstration. The experimental condition that is shown is the one with the HMD and with visual cues enabled. The actual experiment is sped up times eight.

A.2. Unity Project Download

To further improve the reproducibility of this research, the full Unity project can be downloaded at the following url:

<https://github.com/Avdbergnmf/Public-DrossRemoval-Unity-Project.git>

It should be noted that the dross particles have been removed from the project. The reason for this is that the dross particles are simulated using a paid third party asset called Obi Fluid. To be able to fully use the scene, the asset should be purchased, and imported separately from the Unity asset store. The full settings and set-up for the dross particles are discussed in Appendix A.6. However, any other fluid simulator that works in Unity can be used, with slight modifications to the existing code.

The project was created using Unity 2020.1.9f1. All scripts, shaders and models discussed in this appendix can be found in this Unity project. Scripts are located in the *Scripts* folder. The shader for the visual virtual fixtures is located in *Shaders/Custom Shaders/Game Like/Force Fields/Force Field Ripple.shadergraph*. For questions regarding this project, contact the author at: Alex_van_den_berg@hotmail.com.

A.3. External Assets

Most of the environment is custom made and coded. However, several external (third-party) Unity assets have been used. These assets are listed below:

- *Obi Fluid*: A CPU based physics fluid simulator that is well-documented and highly customizable. This package was used to simulate the dross particles. This is the only *paid* package that was used.
- *NVIDIA FleX*: A (GPU based) particle simulator by NVIDIA. This package was used to create a particle based model of the (end-effector) scoop. These particles were then disabled, but their relative positions were used to obtain an approximation of the submerged volume of the scoop. In order to achieve this the package was slightly modified so that the positions of the particles are published, and available from the Unity editor.
- *3DSystems OpenHaptics*: The Unity implementation of the OpenHaptics software, which is the API for the 3DSystems haptic devices. This was used as the basis for the haptic functionality of the scene. On itself, the plugin provides the basic functionality by communicating with the device drivers. This functionality was modified to suit the needs of this project.
- *ROS#*: A package allowing communication between ROS (Robot Operating System) and Unity. This package was used to import a robot model through the interpretation of its URDF's, which contain a description of a that robot.

Additionally, a set of (native) Unity packages have been used. These packages are listed below:

- *Shader Graph*: A visual scripting tool used to create custom Shaders (scripts that inform the GPU on how objects are rendered in the scene).
- *Universal RP*: Unity's scriptable rendering pipeline. This made it possible to greatly improve the frame rate by optimizing the rendering settings.
- *XR Interaction Toolkit*: This package provides a simple way to manage XR (AR/VR) plugins in Unity. Although this asset is still in *preview*, its simplicity allows for easily scriptable functionality, which makes it perfect for this experiment.

A.4. Visual Features

Environment models were provided by a third party to correctly match the target implementation. The materials and their shaders were custom made, although most of the textures are taken from Unity's *measured material library*.

The visual cues are explained quite extensively in the thesis paper. However, the shader for the virtual fixtures is relatively complex and is an interesting topic in itself. Therefore, for the exact details on this shader the reader is referred to Appendix B.

A.4.1. Realistic looking materials

Several shaders were created in order to make the materials as realistic as possible. The idea behind this is that the participant has an accurate perception of the environment and its layout.

The steel strip was made to be quite shiny so that it stands out, as shown in Figure A.1a. It was made to look like the strip is moving by scrolling the textures of the material, as well as animating a small vertex displacement in the shader (scaled by a scrolling gradient noise), making it look like the steel strip is shaking a little.

The shader for the zinc bath surface is special in that it has a real-time reflection of its environment. This is done by placing a second virtual camera behind the surface which is positioned using the position of the used game camera reflected over the bath surface. This image is rendered onto the surface and some effects are applied to mimic the ripples caused by the steel strip continuously running through it. This mirroring effect comes with a lot of additional load on the GPU, as the scene now has to be rendered a second time from another perspective. This gets even worse when using the VR display, as the scene has to be rendered for each eye, and thus, so does the reflection. However, it is vital that this mirror effect is both realtime and fully accurate, as when this is not the case it creates a lot of confusion when the scoop comes close to, or is submerged into, the bath. The resulting look is shown in Figure A.1b.




Figure A.1: Resulting rendering of the custom shaders of (a) The steel strip, (b) the mirroring zinc bath surface resulting from a custom shader.

A.5. Robot Control

The control of the robot through the haptic interface is one of the most important aspects of the scene, and coincidentally one of the most difficult to get right. This is important, as without having a comfortable and intuitive way to control the robot, it will be very difficult to judge the effects of the visual feedback. How this is achieved, is (roughly) described in this section.

A.5.1. Kinematics

The first part is the robot kinematics. Although the robot model itself is imported, its physical properties, and kinematic behaviour is not.

Every link of the robot (other than the base) is made a (physics) rigid body. The colliders of these bodies are simple (convex) mesh colliders (as shown in Figure A.2). The linear drag of all bodies is set to be equal to 2.0, their angular drag is set to 0.05, and their mass is set to 1.0. The correct setting for these parameters is such that unstable behaviour is minimized, yet the robot does not feel *heavy* or *slow*. Additionally, it is important that these drag forces remain low enough so that forces arising from collisions and haptic guidance can still be easily perceived and distinguished from these drag and inertia forces.




Figure A.2: The colliders of the controller robot.

All bodies are linked together using hinge joints. The maximum joint angle of these joints are appropri-

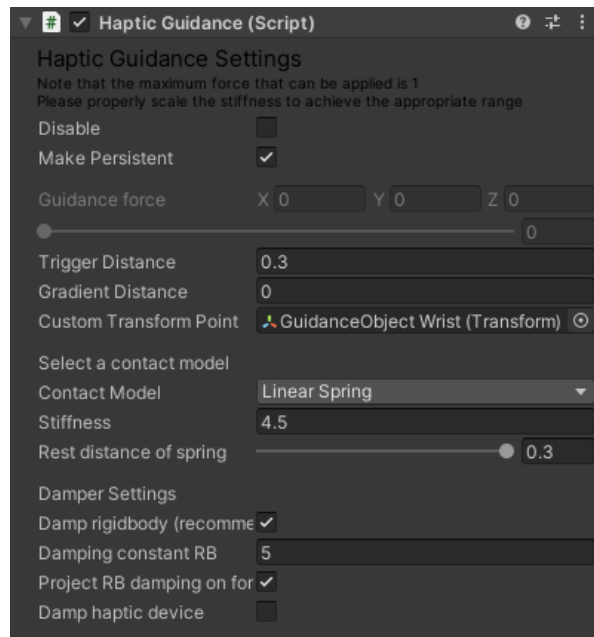


Figure A.3: The editor window of the haptic guidance script.

ately limited to realistically constrain movement past certain angles. This also helps to make the behaviour more stable and intuitive. The joints have a damping value of 2.0. This value is tuned together with that of the rigid body properties stated above.

To mitigate annoying, hard to control behaviour at the singularities, a spring is implemented at the second joint, actively pulling the upper arm (the first link) to an upright pose. The spring constant is 10.0, such that it doesn't impede movement and is barely noticeable, but still functions as intended.

A.5.2. Haptic Control

The haptic control is managed by linking a chosen position (the robot wrist, in this case) to the position of the haptic device using a fixed joint. Due to the set-up of the kinematics, simply moving this point manages the movement of the entire robot.

Movement of the robot results in forces in this fixed joint due to the inertia of the robot, collisions with the environment, and workspace limits. These forces are scaled and damped before being translated to the (real) forces felt in the haptic device. The force is scaled by 0.5 and the damping coefficient is 0.333, once again to ensure stable, comfortable and intuitive control through the haptic interface.

Additionally, it is important to keep in mind the limitations of the haptic interface when tuning these parameters. The used interface can supply a maximum of 2N of force in each direction. If the maximum force is already reached, no additional information can be supplied haptically (in that direction).

A.5.3. Haptic Assistance

The final part of the robot control is the haptic assistance. A custom script was written which applies additional forces directly to the haptic interface. These forces are added on top of the forces that result from the environment interactions, described above. A custom editor script was written to facilitate the fine-tuning of the parameters that define the behaviour of each part of the haptic guidance. The resulting view inside of the Unity editor is shown in Figure A.3.

This script is attached to a game object with a collider. This collider can be of any shape, and defines the area in which the guidance is applied. The forces are activated and scaled based on the distance between the selected transform point (position of a game object attached to the robot) and the collider.

A.6. Dross

A great deal of attention has been paid to the dross particles that are scooped out of the bath. The experiment was designed with the specific usecase of dross removal in mind. In order to get results relevant to this usecase, it is desired that the dross behaves as much as real dross as possible. How this was achieved, is described in this section.

The behaviour of the dross was compared youtube videos of dross skimming¹². The specific settings used for the scripts of the Obi plugins can be found in Figure A.4.

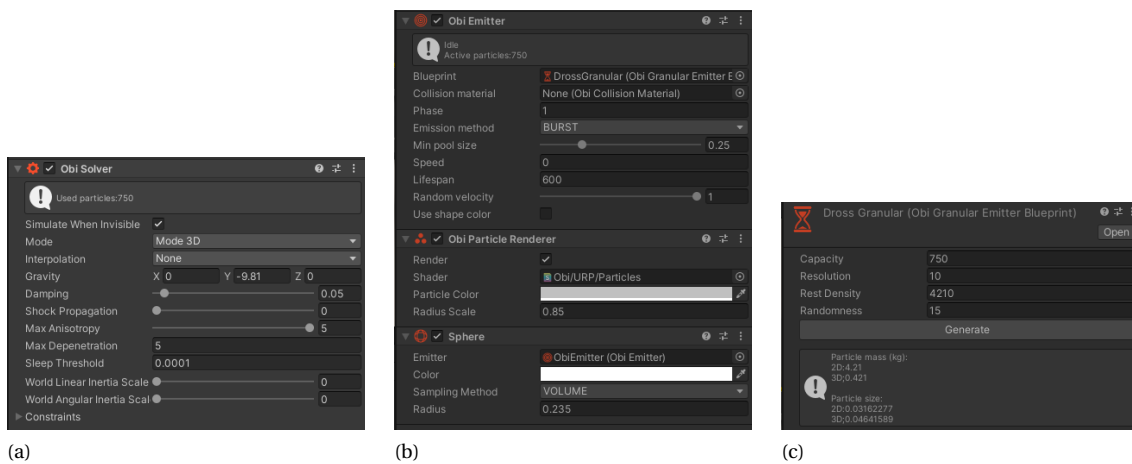


Figure A.4: (a) The editor window of the Obi Solver script, (b) The inspector window of the Obi Emitter game object, (c) The editor window of the Obi Blueprint for the dross particles.

Additionally, the way the dross particles interact with the environment also greatly affects the behaviour of the particles. This interaction is defined by assigning Obi materials to the colliders with which the particles will interact. Most importantly, the friction of the zinc bath surface is set to be relatively low, while that of the other materials is relatively high. Additionally, the scoop is made slightly sticky. Because of this, the particles are more stable when transported in the scoop and don't fall out as easily, while still being easy to dump in the bin. The values of all parameters for all materials used are shown in Figure A.5.

¹DROSS SKIMMING ROBOT - ABB

Mustafa Neiyi

https://www.youtube.com/watch?v=1skq__LMC_M

²DROSS SKIMMING ROBOT IV / Robomax Robotic

Robomax Robot ve Otomasyon Sistemleri

<https://www.youtube.com/watch?v=1LHwTKU1PTk>




Figure A.5: The editor window of the Obi material scripts. These are all the materials that were used in the scene.

A.7. Logging and Data processing

Another important aspect of the Unity project is the logging of the data. Unity has no default loggers, but with a set of C# scripts everything can be logged to csv files quite easily.

For each type of data recorded, a separate script was written, all of which derive from the same class called *logToCSV*. This class contains the basic functionality such as creating the log file and writing to it. All log scripts are then linked to a logging manager script. This script makes sure all logs are created into the correct folders, and start and stop logging at the right times.

The data was later analysed using Python scripts, using the following libraries:

- *Matplotlib* - Used for data visualization
- *Seaborn* - Used for data visualization
- *Pandas* - Used for processing the data
- *SciPy* - Used for the statistical analyses

The full code used for the analysis can be downloaded from the following url:

https://github.com/Avdbergnmf/DataProcessing_DrossRemovalThesis.git

B

Virtual Fixtures Shader

The visualization of the virtual fixtures (VFs) is an important part of this research. Because the way in which this visualization is set up is relatively complex, the full details are explained in this appendix.

The visual VFs were created using a specially written shader. Shaders are used to create materials inside of Unity, which determine how an object is rendered on the screen. The shader for this visualization is created using Shader Graph, which is a visual shader scripting method by Unity. This shader graph will be used to explain how this shader works. Although the shader can be compiled into code, this will likely not provide much insight for the average reader. Instead, the shader graph is decomposed into several parts, each of which is explained in this appendix. Theoretically, the reader can recreate this shader themselves, or base a new one on the ideas presented here.

Parameters that are material specific (called *properties*) need to be edited from the Unity editor. These will be signified by being printed in *italic* in this appendix. The values of these properties are shown in Figure B.1. Note that the *Emission* defines the color of the material, which is different for the three different types of virtual fixtures.

The full graph is shown in Figure B.2, showing a complete overview of its components. The texturing part (A) is shown in Figure B.3. This implements a simple hexagonal repeated *Pattern* (texture), which scrolls over the object with a speed of *Scroll Speed*. This texture is then fed into the alpha (transparency) of the material making it visible. Without a texture it is a lot more difficult to understand the shape and other spatial features of the visualized virtual fixtures.

The sphere mask (B) of the shader is shown in Figure B.4. A sphere mask (of radius *Sphere Radius*) is implemented, inside of which the alpha is increased (locally making the material brighter/more opaque). The *Sphere Center* is linked to the location of the points used to create the haptic guidance (using a C# script). As a result, the material gets brighter around these linked locations.

The boundary circle (C) is the part of the shader that is responsible for the outer ring of the brighter area defined in (B). The boundary circle works quite similar to B, also using a sphere mask (with radius *Sphere Radius*). However, the position of the sphere mask *Boundary Circle Center* is determined by the location on the virtual fixture's collider that is closest to the *Sphere Center* of B (using another C# script). In this way, the boundary circle always has the same radius, and is centred around the bright spot caused by (B). A zoomed in screenshot of this part of the shader is shown in Figure B.5.

When the location of the brighter sphere (*Sphere Center*) is inside of the collider (*Circle Center* is equal to *Boundary Circle Center*) the brighter area will start blinking. This behaviour is defined in (D), shown in Figure B.6. Here, *Enable Blinking* is changed to True to enable the blinking.

Finally, the components are combined in part (E) as shown in Figure B.7. The Balance part on the left simply applies some scaling to make sure the radii of the sphere masks are always correct. The right part shows how the components are combined and fed into the material to create the end-result.




Figure B.1: The properties and their values of the virtual fixture shader.

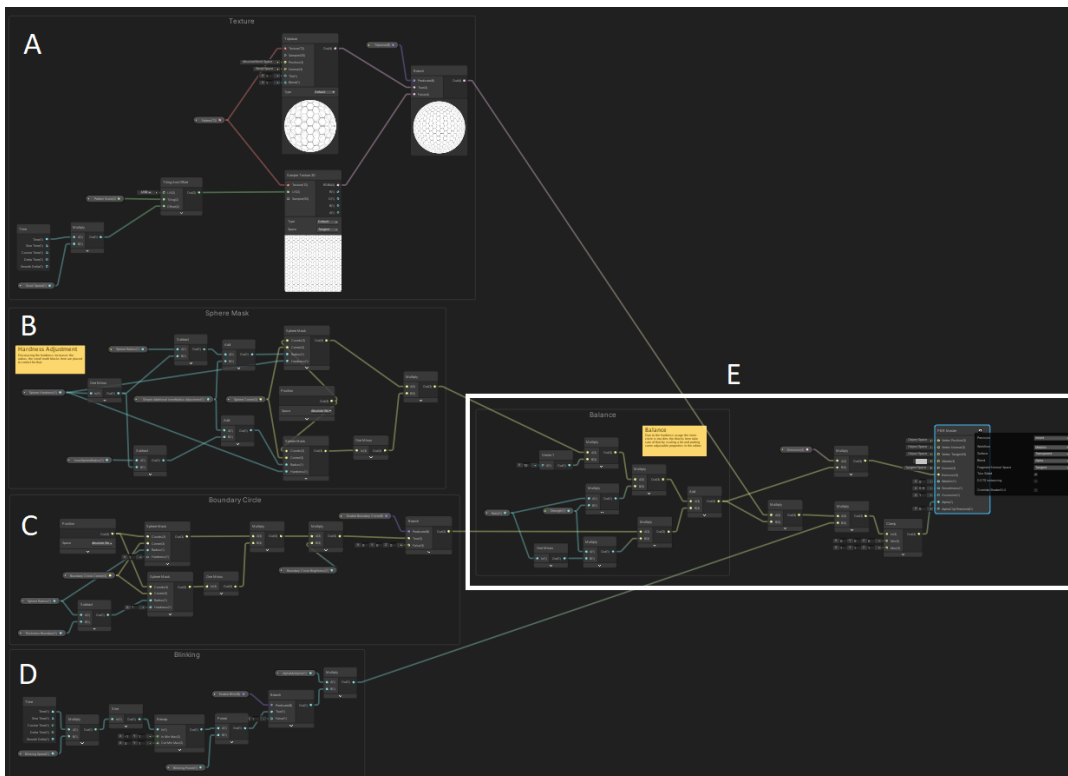


Figure B.2: A screenshot of the full shader graph, showing the overall layout of the separate components and how they are combined.




Figure B.3: Part A: Texture.

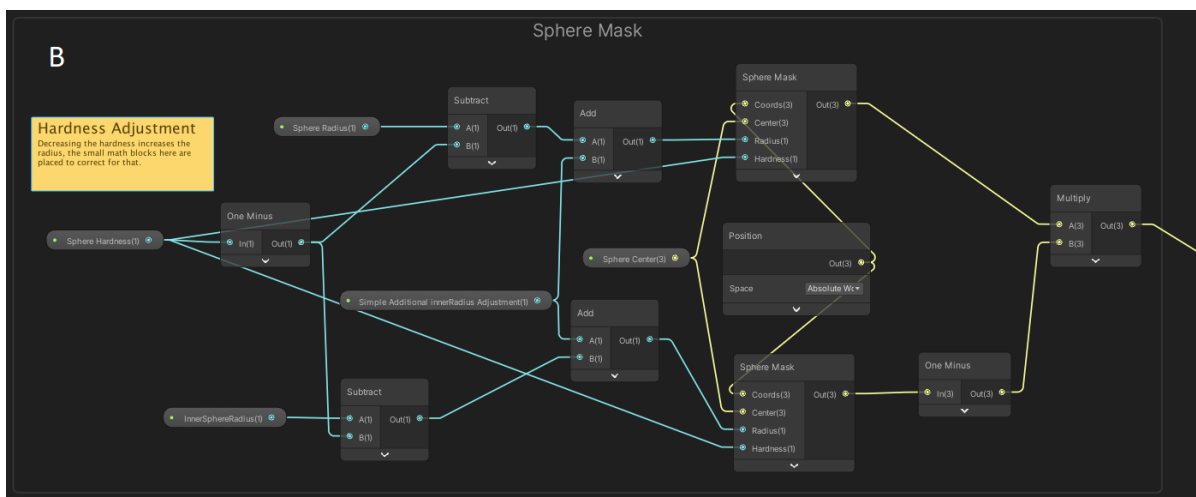


Figure B.4: Part B: Sphere Mask.

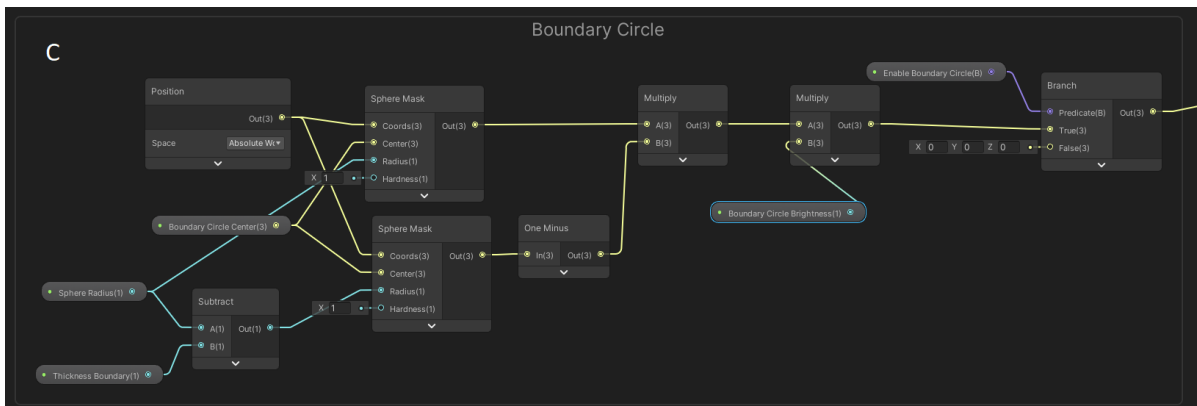


Figure B.5: Part C: Boundary Circle.




Figure B.6: Part D: Blinking.




Figure B.7: Part E: Combining the components.

C

Participant Instructions

Participant information before starting – Goal, method and setup of the experiment

Introduction

You are remote operator of a **simulated** industrial robot (see Figure 4).

- Your job is to remove dross (a solid contamination in the form of small balls) from a liquid (molten) zinc bath.
- This zinc bath is used in a galvanization process of a steel strip continuously running through this bath at high speed. The layout of this environment is shown and explained in Figure 1.
- There are a couple of important aspects, that you as the operator in this simulation need to take into account:
 - Touching the steel strip will cause a lot of (financial and physical) damage – doing this in the simulation during the experiment will cause the trial to **fail** instantly
 - Touch the environment, especially with high impact forces, will cause damage to both the environment and the robot – Exceeding a set collision force in the simulation during the experiment will cause the trial to **fail** instantly
 - Causing disturbances in the bath deteriorates the quality of galvanization achieved by the galvanizing process




Figure 1 - The simulation environment. Here 1 is the **main menu** for the experiment, which will be shown at the start of every experiment and shows the currently active cues. 2 indicates the **controlled industrial robot**, also shown in Figure 4. Number 3 indicates the **dross collection bin**, in which the dross should be deposited. 4 indicates the **zinc bath** with the **dross (balls)** floating on top. Number 5 indicates the **steel strip** which is undergoing the galvanizing process. In Figure 2 a topdown view is shown of this same environment, here the number 6 indicates the **approximate position from which the robot will be viewed**, which is the point of view in this figure.

Goal (of the experiment)

- To aid in this task, this experiment aims to evaluate if, and the degree to which visual feedback contributes to the performance, and operator experience in performing this task.




Figure 2 - A topdown view of the simulated environment also shown in Figure 1, where the meanings of the numbers are explained in detail.

Method

- To achieve this, a simulation environment is set up in which the performance of this task can be simulated. There are four experimental conditions, which will be presented to you in a predetermined order:
 - Displayed on a monitor, with all visual cues disabled
 - Displayed on a monitor, with all visual cues enabled
 - Displayed on a VR head-mounted display, with all visual cues enabled
 - Displayed on a VR head-mounted display, with all visual cues disabled
- Every visual and haptic cue will be explained prior to the experiment, and you will be requested to test out, understand and get a feel for each of them during this explanation.
- The commanding of the robot will be achieved by interacting with the 3Dsystems Touch, a haptic, 6 degree of freedom telecommunication “joystick”, that transfers haptic feedback in translation to the teleoperator (see Figure 3). You can toggle the connection to this robot by pressing the dark button (button 2). You will get a chance to test this out and get a feel for this before starting the experiment.



Figure 3 - The Phantom Omni (<https://www.3dsystems.com/haptics-devices/touch>). The stylus is the pen-like device inserted into the device when not used. The 1 and 2 indicate the buttons that are located on this stylus. Button 1 is the one on the back (white/light) and button 2 is the one on the front (black/dark).




Figure 4 - The industrial robot operated during the experiment. The most important parts of the robot are indicated with captions inside the figure. The scoop is the object that will be used to scoop out dross from the bath. It is attached to the wrist of the robot, which is going to be the controlled point. The base is stationary and the middle of the workrange of the robot.

Evaluation Criteria

Your performance in this experiment will be evaluated on the following criteria:

- Success – did you manage to remove the dross without touching the steel strip or causing collisions with a high impact force?
- Safety – number of collisions and the maximum impact force
- Speed - How much dross are you able to scoop out?
- Accuracy – Did you accidentally drop some dross outside of the collection bin?
- Bath Disturbance – To minimize disturbance to the bath, it is important you try to not submerge the scoop more than necessary and to not move inside the bath faster than necessary.

Experiment

Here, a thorough overview of the course of the experiment is given. You will be guided through this process each step of the way, though it will help to give you a better understanding of what you can expect.

- Experiment set-up:
 - After receiving all necessary explanation, you (the operator) will be presented with the described simulation environment of the zinc bath with a predetermined set experimental condition. You are able to see the state of this condition before starting every trial.
- Experiment overview:
 - For each of the conditions you will start off by performing a practice trial to familiarize yourself with the current set up.
 - This practice trial is not evaluated and serves to get the operator used to the current condition set. However, it is important that you try your best to complete it to the best of your ability.
 - You will perform 3 repetitions for each condition, lasting 5 minutes per repetition.
 - After each trial you will take a short break.
 - After completing all trials for a condition, you will be asked fill out a questionnaire in which your are asked to evaluate the experience.
- Trial overview:
 - In each trial your objective is the same: You will remove as much of the dross as you can (represented by small spheres/balls) from the zinc bath and deposit them in the dross collection bin (see Figure 1).
You will do this as carefully, and safely as you can, while still trying to do this in a timely manner.
 - The trial will end under one of these conditions:
 - The time frame of 5 minutes has passed.
 - You decide to withdraw from the trial/experiment (there will be no negative consequences for you as a result of this, as described in the consent form)
 - You fail the experiment by either:
 - Touching the steel strip with the robot
 - Colliding the robot with the environment with a high impact force

Participant Consent Information

Visual feedback for haptic teleoperation with dress removal as a use case

Please read this information and all other information provided to you carefully, including the “Participant information before starting”, and ask any questions to the researcher if you have them.

Withdrawing participation

Note that, at any time, you can withdraw from participation of the experiment without having to give a reason and without any punishment for doing so.

Personal Information

For this experiment your age and gender will be asked at the start. This information is collected to that trends relating age/gender to performance might be observed. This information is stored anonymously, and is not linked to your name or any other personal information. The data will be stored according to the [TU Delft Research Data Framework Policy \(https://www.tudelft.nl/en/library/current-topics/research-data-management/\)](https://www.tudelft.nl/en/library/current-topics/research-data-management/). The policy stipulates that research data should be retained for at least 10 years from the end of the research project, unless there are valid reasons not to do so.



The data is stored on the TU Delft research data repository (accessible by project members: the researcher, Alex van den Berg), and on the Heemskerk Innovative Technology (HIT) company repository (accessible only by HIT personnel).

Contact

If you have any questions about the research you can contact the appropriate person by using the contact information below.



Contact details:

- Researcher: Alex van den Berg, A.vandenBerg-2@student.tudelft.nl, +31642149934
- Data Protection Officer: Drs. Y. Türkyilmaz-van der Velden, Y.Turkyilmaz-vanderVelden@tudelft.nl
- University Supervisor: Prof.dr.ir. D.A. (David) Abbink, D.A.Abbink@tudelft.nl
- Company Supervisor: Cock J. M. Heemskerk, PhD, c.heemskerk@heemskerk-innovative.nl

VR Headset – VR sickness

For part of the experiment you will be wearing a VR headset. Although the experiment is set-up in a way to minimize this, there is a risk of experiencing VR sickness, resulting in nausea and 'dizziness'. If you experience any type of discomfort during this experiment, please report this to the experimenter immediately so that the experiment can be halted. If you are known to get motion-sick quite easily, please indicate this to the experimenter so that he is aware of this.



COVID-19 precautions

During this experiment you will come into physical contact with some digital devices (the Vive pro VR display and the 3DSystems Touch, haptic joystick). This comes with the risk of being exposed to COVID-19 (the coronavirus). To minimize this risk, the researcher and participant (you) ensure the following precautions:



- Contact time between experimenter and participant (you) is minimized by sending information prior to the experiment if possible.
- Any object the you might come into contact with for this experiment is thoroughly cleaned using soap or alcohol wipes in between every participant
- A special room is set up where the experiment takes place. In this room at most only the experimenter and the participant are present.
- A distance of 1.5 meter between the participant and experimenter is maintained throughout the entirety of the experiment (social distancing).
- You are asked to wash your hands when arriving at the experiment location, and before leaving.
- **If you are experiencing symptoms of a cold/flu, or live with people who do, please stay at home and refrain from participation.**
- **If you need to travel by public transport to attend this experiment, you are asked not to, and refrain from participation.**

D

Participant Consent Form

Consent Form for Visual feedback for haptic teleoperation with dross removal as a use case

Please tick the appropriate boxes

Yes No

Taking part in the study

I have read and understood the study information dated 20/04/2020, or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.

I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.

I understand that taking part in the study involves the capturing of spatial and force data in a simulated environment, and a survey questionnaire completed by the you (the participant)

Risks associated with participating in the study

I understand that taking part in the study involves the following risks:

- mental discomfort due to VR motion sickness
- An increased risk of exposure to Covid-19 (the coronavirus)

Use of the information in the study

I understand that information I provide will be used for a possible paper publication, the resulting report for the researchers master thesis, the possible development and/or adjustment of a control system of a dross removal system (the target use-case of this study), possible publication of the (aggregated) experiment results on the company website (<https://heemskerk-innovative.nl/>)

I understand that personal information collected about me that can identify me, such as age and gender, will not be shared beyond the study team.

Future use and reuse of the information by others

I give permission for the anonymised logged simulation data (positional, and force data) and questionnaire results along with my age and gender that I provide to be archived (in aggregated form) in TU Delft Research data repository so it can be used for future research and learning.

Signatures

Name of participant

Signature

Date

I have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

Researcher name

Signature

Date

For further information, contact the researcher:

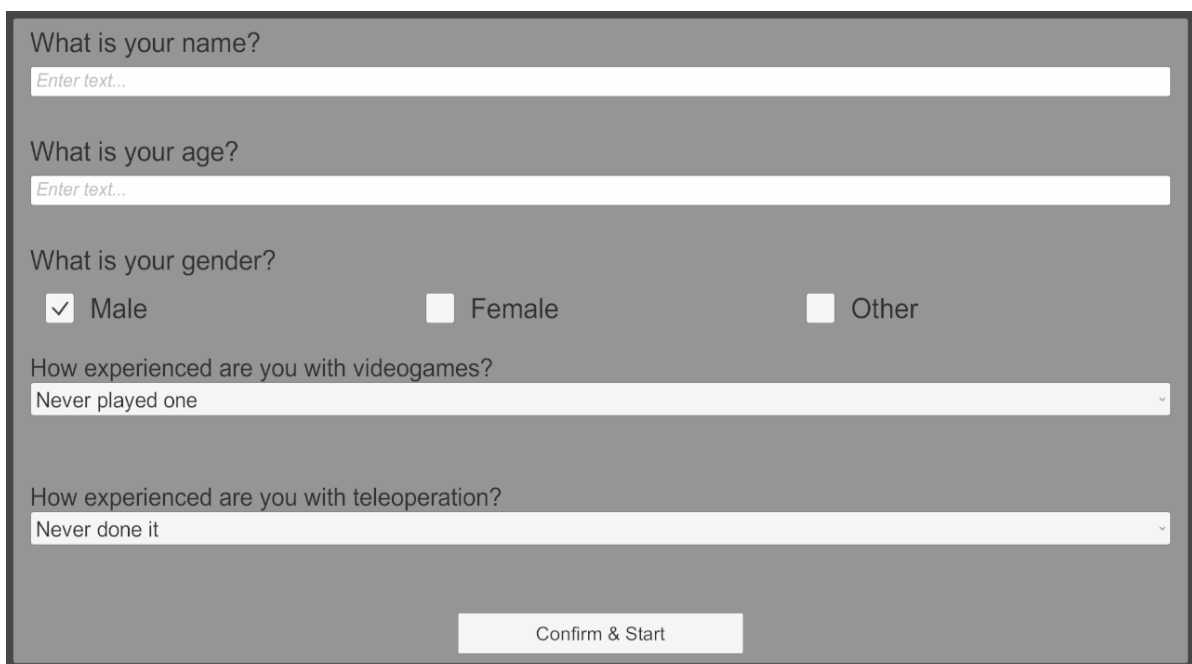
Alex van den Berg, +31642149934, Alex_van_den_berg@hotmail.com

E

Questionnaires

This appendix shows screenshots of all the questionnaires. These questionnaires were filled out and logged inside of the Unity environment.

This is the starting questionnaire, in which participants were asked to fill out some basic info that was used for the demographics.



What is your name?

What is your age?

What is your gender?
 Male Female Other

How experienced are you with videogames?

How experienced are you with teleoperation?

This screenshot shows the van der Laan[1] user acceptance questionnaire. The sliders could only be placed at integer values.

My judgements of the current control interface (as a whole) are ...
(please appropriately adjust the slider on every line)

1. Useful	<input type="range"/>	Useless
2. Pleasant	<input type="range"/>	Unpleasant
3. Bad	<input type="range"/>	Good
4. Nice	<input type="range"/>	Annoying
5. Effective	<input type="range"/>	Superfluous
6. Irritating	<input type="range"/>	Likeable
7. Assisting	<input type="range"/>	Worthless
8. Undesirable	<input type="range"/>	Desirable
9. Raising Alertness	<input type="range"/>	Sleep-inducing

This is the final questionnaire. The comments written in these answers were used to gain additional insight into the results of the dependent measures.

Please answer these final few questions:

What, if anything, did you think was difficult about this experiment? Please try your best to provide examples when you can.

Which condition did you like the best?

1	<input type="checkbox"/>	VR + No Cues	2	<input type="checkbox"/>	Monitor + Cues	3	<input type="checkbox"/>	VR + Cues	4	<input type="checkbox"/>	Monitor + No Cues
---	--------------------------	--------------	---	--------------------------	----------------	---	--------------------------	-----------	---	--------------------------	-------------------

Why did you like this condition the best? Please try your best to provide examples when you can.

If you have any additional remarks, please type them in here...

F

Shapiro Wilk Test of Normality Results

The Shapiro Wilk results for all datasets are shown in Table F.1. The abbreviations for the experimental conditions are as follows:

- *MN*: Monitor display, and No visual cues
- *MC*: Monitor display, and with visual Cues
- *HN*: HMD, and No visual cues
- *HC*: HMD, and with visual Cues

Table F.1: Resulting values of the Shapiro Wilk test of normality. Significant values ($p \leq 0.05$) indicate non-normality and are printed in **bold**.

dataset	MC	HC	MN	HN
w value: Number of scoops	0.977	0.975	0.967	0.961
p value: Number of scoops	0.932	0.906	0.787	0.688
w value: Average scoops size	0.943	0.919	0.943	0.967
p value: Average scoops size	0.385	0.164	0.386	0.794
w value: van der laan Usefulness	0.911	0.818	0.808	0.935
p value: van der laan Usefulness	0.123	0.005	0.004	0.295
w value: van der laan Satisfying	0.914	0.843	0.917	0.932
p value: van der laan Satisfying	0.135	0.011	0.153	0.265
w value: Maximum scoop flux	0.897	0.763	0.886	0.912
p value: Maximum scoop flux	0.072	0.001	0.048	0.126
w value: Minimum distance to steel strip	0.889	0.827	0.903	0.88
p value: Minimum distance to steel strip	0.054	0.006	0.09	0.038
w value: Dross Removed	0.977	0.962	0.926	0.961
p value: Dross Removed	0.939	0.703	0.21	0.685
w value: Average Disturbance	0.846	0.693	0.714	0.619
p value: Average Disturbance	0.012	0.001	0.001	0.001
w value: Average velocity	0.969	0.967	0.955	0.959
p value: Average velocity	0.827	0.793	0.575	0.651
w value: Maximum collision force	0.949	0.922	0.962	0.919
p value: Maximum collision force	0.474	0.181	0.693	0.161
w value: Accuracy dross dumping	0.9	0.944	0.802	0.812
p value: Accuracy dross dumping	0.079	0.398	0.003	0.004

G

Two-way RM ANOVA Results

Table G.1 contains the results of the two way repeated measures ANOVA for all recorded and calculated data. Data that did not pass the Shapiro Wilk test of normality F was transformed using the Aligned Rank Transform [2].

Table G.1: Results of the two way repeated measures ANOVA for all recorded and calculated data. Significant effects ($p \leq 0.05$) are printed in **bold**.

dataset	Cues: Pr >F	Cues: F Value	Display: Pr >F	Display: F Value	Cues:Display: Pr >F	Cues:Display: F Value
Number of scoops	0.63148496	0.239729	0.117467	2.758908	0.828601	0.048536
Average scoopsize	0.93171451	0.007593	9.76E-05	27.58361	0.531416	0.410432
van der laan Usefulness	0.00120623	15.8444	0.000959	16.75566	0.040526	5.025379
van der laan Satisfying	0.010018508	8.677755	9.56E-06	42.6087	0.131645	2.54284
Maximum scoop flux	0.019980333	6.775476	0.059458	4.158173	0.896845	0.017388
Minimum distance to steel strip	0.139897202	2.429752	0.351844	0.9233	0.53415	0.404919
Dross Removed	0.786007849	0.076401	0.003874	11.63113	0.715412	0.138062
Average Disturbance	0.088649156	3.315181	0.857337	0.033447	0.398865	0.75409
Average velocity	0.214134533	1.682927	0.037148	5.230168	0.868701	0.028279
Maximum collision force	0.000362958	20.95179	0.669041	0.190102	0.313844	1.086096
Accuracy dross dumping	0.306780221	1.119444	0.689203	0.166277	0.967566	0.001709

H

Other Results

The boxplots and barplots of all the results that were recorded and calculated are displayed in this appendix. Dependent measures that were not discussed in the scientific paper are briefly explained in the figure caption. The marks denote significance, where • $p \leq 0.05$, •• $p \leq 0.01$, ••• $p \leq 0.001$.




Figure H.1: (a) The amount of dross that landed in the bins divided by the amount of dross that was taken out of the bath (b) Average Disturbance (c) Average Scoop Size




Figure H.2: (a) Average Velocity (b) Fraction of Dross removed (c) Peak Collision Force




Figure H.3: (a) Maximum value for the change in the submerged volume (fraction removed difference/stepsize) (b) Minimum Distance to Steel Strip (c) Number of Scoops




Figure H.4: (a) van der Laan Satisfying Scores (b) van der Laan Usefulness Scores

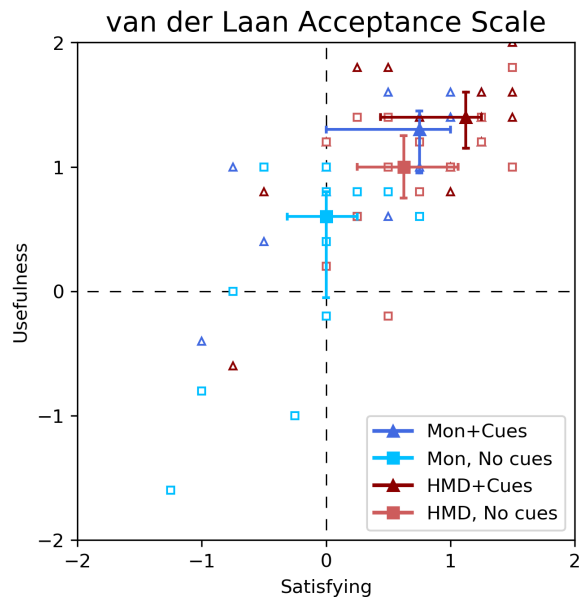


Figure H.5: Van der Laan acceptance scores[1]; on the horizontal axis the satisfying dimension and on the vertical axis the usefulness dimension. Scores are in the range of (-2,2). The center points are the median values, and the colored lines show 25% percentile ranges. The other marks are the separate datapoints of all participants.




Figure H.6: (a) Number of Failed Trials (b) Preferred Condition

Bibliography

- [1] Jinke D Van Der Laan, Adriaan Heino, and Dick De Waard. A simple procedure for the assessment of acceptance of advanced transport telematics. *Transportation Research Part C: Emerging Technologies*, 5 (1):1–10, 1997.
- [2] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 143–146, 2011.