Now You’re Teleoperating with Power: learning from video games to improve teleoperation interfaces

By

Daniel J. Rea

A Thesis submitted to

the Faculty of Graduate Studies of the University of Manitoba

in partial fulfilment of the requirements of the degree of

Doctor of Philosophy

Department of Computer Science

University of Manitoba

Winnipeg

Copyright © 2020 by Daniel J. Rea

ABSTRACT

Teleoperation has potential applications in the home, industry, and other areas such as search and rescue. Safe and efficient teleoperation is difficult, however, and improved interaction design is one way to mitigate the challenges operators encounter. Video games share many similar challenges to teleoperation in terms of interaction design: both have a user controlling an entity in a remote space, receiving feedback and sending controls. I investigate how to improve teleoperation performance and experience by learning from video game interaction design.

For years, video game developers have been creating interactions in their games (e.g., in game events, interface elements, and characters) that influence how their players think and act by leveraging different aspects of human psychology. I investigate how I can take design cues and inspiration from these psychology-based video game interaction to design, implement, and evaluate new interaction designs that consider or shape how operators think and act during teleoperation.

I successfully design and experimentally evaluate a concrete set of video game-inspired teleoperation techniques from three perspectives: directing operator attention, priming operator perceptions of robot capability to shape driving behaviour, and using social agents to influence operator experience and driving behaviour. In addition, I create a framework of video game interaction design; my framework provides the structure and vocabulary for discussing video game interactions at an abstract level, which I leverage to showcase the similarities of the problem spaces, and solutions.

DEDICATION

To Mom, Grandma, and Papa:

You didn't get to see me finish, but you always believed in me,

And I know you would be proud.

ACKNOWLEDGEMENTS

While this Ph.D. may seem to be the work of myself alone, in reality I could only get to where I am thanks to the support and help from so many people.

I would like to start with thanking my committee. Throughout the years they provided discussion, new ideas, insights, critical commentary, and guided me to make this thesis something I am proud of.

A special thank you to my supervisor, Jim. It’s been a long ride of many many (many) years. Thank you for the tough talks, the mentorship, the opportunities, and yes, even the seemingly endless edits you would put into each of my drafts. I honestly am not sure there is a better supervisor out there, and I would do it all over with you again. You have prepared me to be a professional in research, in teaching, and as a member of a community, but you also did it while being a friend. I am proud to be your first Ph.D. student.

To my family, particularly Dad and Megan, who supported me through the exhausted weekends, the flus, the endless rants, the panics, and of course, the successes. Losing so many important family members during this Ph.D. was harder than the degree itself. Thank you for being there throughout this journey. This is also your achievement. I love you!

Just as important, to the many many friends who helped, read drafts, grabbed beers, drank coffees, watched presentations, ripped presentations apart, helped me relax and remember

there was a life outside of research, shared in my successes, and were compassionate during hard times. I treasure all the days and nights in the lab with my labmates, all the adventures with conference buddies, all the coffees with old friends, all the movie nights, the coding jams, the swing dances, the bike rides, the gym sessions, the walks, the talks, and the therapy gaming. I may be flying around the world doing research now, and maybe we can’t see each other as often, but you are all always in my heart. You being there for me throughout these years, even if just for the rare meet up and chat, means the world to me.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter One:

An Introduction to the Similarities between Teleoperation and Video Games

1.1 Thesis Overview

Teleoperated robots, or robots that a human operator controls from a distance, are an increasingly important part of many industries. Applications of teleoperated robots have promised a future where we can save lives with remote search and rescue robots (Borenstein and Pearson 2011), improve clinical care (M. a. Goodrich, Crandall, and Barakova 2013), conveniently control robots in our homes (Liu et al. 2011), improve customer service (Glas et al. 2012; M. a. Goodrich, Crandall, and Barakova 2013), and more. The primary benefit in all these cases is enabling a person to act in another, perhaps remote, physical space; with teleoperation, robot operators can enter dangerous areas to work such as cleaning up nuclear disasters (Manocha, Pernalete, and Dubey 2001), apply highly specialized skills such as medical knowledge to communities that lack such specialists (Cai Meng et al. 2004), or simply attend a work meeting on a day they have to stay home waiting for a repairperson. As robotic technology and wireless networking continues to improve, the applications and accessibility of teleoperation will continue to improve numerous jobs and lives.

Teleoperated robots, however, are very difficult to control (Drury, Scholtz, and Yanco 2003; Yanco and Drury 2004a). One major reason is the difficulty of understanding and monitoring the remote robot and its environment (Endsley 2016; Endsley 1988) due to people not being able to leverage their normal senses (wide field of vision, spatial sound localization, understanding our body posture, and more). In addition, robot operators may have complex objectives for their robots in dynamic and unpredictable environments (Taylor, Yanco, and Drury 2009). For example, a robot may be moving through a crowd of people and a pedestrian may suddenly cross in front of it, requiring quick operator reactions to not hurt anyone. Robots may also be used to put out fires or help in other emergencies, where the operator may need to suddenly change goals while trying to react safely and quickly to the current situation. These circumstances require an operator to build and maintain knowledge of the environment around the robot and to be able to quickly create and adapt strategies to new information, resulting in high cognitive load during operation. Because of this issue of building and understanding complex data from the robot and its surroundings, many problems in teleoperation are linked to human perception and information processing, and are known to be a significant cause of mistakes and incidents during teleoperation (Giese, Carr, and Chahl 2013; Williams 2004; Cotter 2014).

Thus, teleoperation presents many challenges to the operator: robots have various capabilities to control, and operators must navigate robots through complex environments such as disaster sites or offices and complete goals (attend meetings, search and rescue, etc.) with a limited understanding of the remote environment. I note that many of these monitoring and control problems are similar to interaction challenges present in many video games. For example, similar to controlling a complex robot, video game players control an avatar with a range of functions and abilities; like using a robot to explore and search a crowded mall or disaster scenario, game players must navigate an avatar (or vehicle, etc.) through a dynamic environment to achieve objectives; all the while, just as with robot operators, game players must build and maintain an awareness of their avatar, the avatar's status in the game and the virtual environment, relying only on limited visual, sound, and haptic feedback. This is true for many styles or genres of games including first-person shooters, racing games, or role-playing games. Due to the similarities of the problems between the two fields (Figure 1.1), I argue that the interface techniques that have emerged to improve video game interaction and mitigate related problems can and should be used for improving teleoperation interfaces.

We note that the video game industry has dedicated decades of work to addressing this problem of remotely controlling an avatar (e.g., Burgess 2014; Candland 2016; D. Sakamoto 2015), work that we can learn from in developing interfaces to remotely control robots. Further, video games are now ubiquitous in the general population, where half of all Americans, for example, play several hours per week, with a near equal gender balance (*Essential Facts about the*

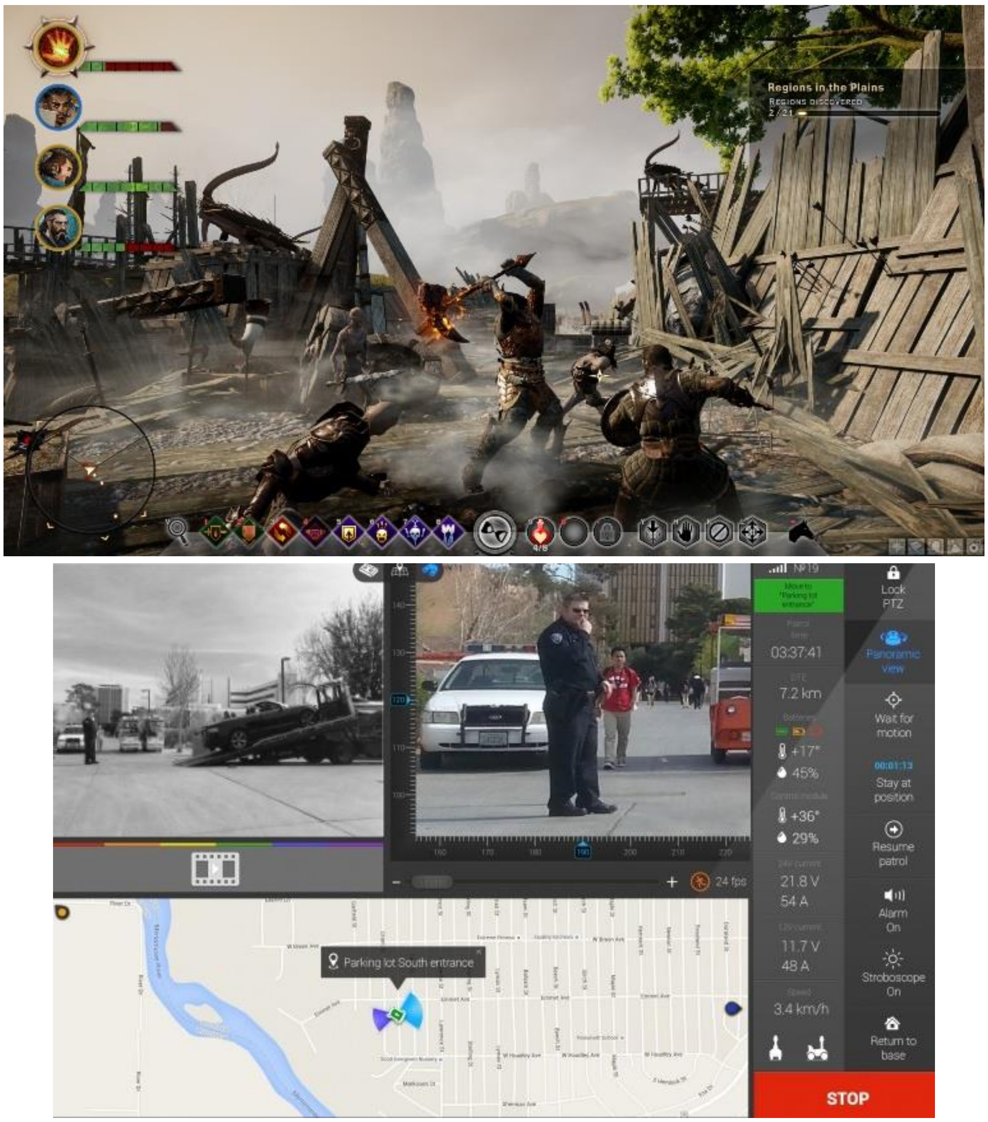

Figure 1.1 Video games and telerobotics interfaces share similar interaction design problems. Both the top [Dragon Age: Inquisition, BioWare, 2014] [1] *and the bottom teleoperation interfaces (SMP Robotics) display video feeds, navigational aids, status visualizations, available commands, etc. These must be designed in an understandable and usable manner.*

[1] In this thesis, we maintain a separate list of references for video games. Citations appear in square brackets, include the game title, the developing studio(s), and year of publication. A full list of the video game references we use are presented in Appendix D.

Computer and Video Game Industry 2015; Grundberg and Hansegard 2014; Warman 2011). Therefore, in addition to learning from video game interfaces for solutions to problems shared with teleoperation, video game design elements in teleoperation interfaces can leverage peoples' experience and knowledge of video games, decreasing the need for training. In this thesis I systematically investigate this proposal – learning from video game interfaces to improve teleoperation – by exploring three specific relevant interaction challenges and presenting a framework of video game interaction techniques focused on how they can be applied to teleoperation.

Using common video game interface techniques as-is in robot teleoperation, however, poses several challenges; video games operate in virtual worlds that afford designers freedoms not possible in the real world. For example, a game interface designer can know where and what everything is in the environment with near-perfect accuracy in real time, and are able to modify game reality and physics –such as ignoring damage from falling, or having a view camera that can move freely and even see through walls to improve the player's view. In contrast, telerobotics must often deal with difficult conditions such as noisy sensors, high latency networks, and an unknown world with potential dangers (e.g., in robot urban search and rescue). Thus, despite similarities in terms of problems faced, video game interfaces may not be trivially applied to teleoperation, and requires research to verify if and how these interfaces can benefit teleoperation.

A key element of video game design has been its focus on users. Video game design has engaged with a range of issues surrounding human limitations: people have limited attention (Pylyshyn and Storm 1988), they misinterpret information (Recarte and Nunes 1996), are unconsciously influenced by their environment (Aarts and Dijksterhuis 2003; Ackerman, Nocera, and Bargh 2010), and have many more innate tendencies, reflexes, and ways of thinking that can be studied and leveraged by good design. As games are virtual, designers can manipulate interfaces and environments to use these tendencies of human behaviour and psychology to impact how players interact with the game and how they feel about those interactions (Langer, Hancock, and Scott 2015; Scheurle 2017).

Teleoperation interface research already employs knowledge of human factors and psychological effects to improve interaction as an ongoing research theme. Researchers have been creating interfaces and control schemes that reduce cognitive load during teleoperation to help operators spend less time and energy on control and interpreting sensor data, and to enable operators to better focus on higher level goals (e.g., D. Sakamoto et al. 2009; Hashimoto et al. 2011; Saakes et al. 2013). In this thesis, I continue this approach and focus on studying how we can leverage video game interaction design techniques that shape user experience to similarly shape teleoperation.

1.1.1 Research Objectives

Throughout my research, I continually investigate the following broad research questions:

1) What types of interfaces exist in video games and how would they be beneficial to teleoperation?
2) How do human factors or psychological mechanisms used in video games affect teleoperation?
3) How do game user experience goals change the teleoperation user experience?

To investigate these objectives, I follow a human-centred approach, aiming to learn from human abilities and psychology to shape the user experience and impact user behaviour. I design novel teleoperation interface techniques, learning from video games, and evaluate operator experience and their teleoperation performance. Specifically, I investigate three angles for how I can learn from video game interface design, and look to the broader applicability of my approach by developing a framework of video game interaction:

a. Human Perception and Attention: video games leverage knowledge of low-level human perception to draw attention, manage distraction, and keep a game engaging. I draw inspiration from video-game visual interface techniques and knowledge from the perceptual psychology literature to investigate how to effectively draw the attention of robot operators to important areas in a robot video feed with visual interfaces.

b. Priming Operator Expectations. Video games orchestrate user expectations and experience using design techniques, such as how a character or vehicle is presented in-game, or how such a vehicle responds to a user command. In order to take advantage of similar effects in teleoperation, I build from these video game techniques to explore how to prime people's expectations and perceptions of a robot's abilities. I then observe how that affects the operator's teleoperation behaviour and experience.

c. Social Interfaces in Teleoperation. Video games often have virtual non-player characters (in addition to the player's character) that react to the player's in-game actions such as by being emotionally hurt or using persuasive dialogue; players may react to these social actions as if they were performed by real people, changing the player experience and even the player's behaviour. I investigate the effects on the operator of incorporating social agents into the teleoperation interface to elicit similar effects in teleoperation.

d. Video Game Interaction Design Framework. To consider the broader applicability of my strategy, I finish this thesis with a high-level analysis of the video game interface landscape from the perspective of teleoperation. Specifically, I provide a multi-layered framework for describing the techniques used in video games, including the user experience goals, high-level implementation strategies, specific implementation choices, and the types of information sent between the player and game. My framework helps facilitate future work in this field by providing researchers and practitioners with tools to discuss video game interfaces and the interaction problems they target in terms of their applicability to teleoperation; I use my framework in this way to highlight the similarities of teleoperation and video game interaction design.

The first three projects constitute proof-of-concepts for my approach of learning from video games for designing teleoperation interfaces. My framework serves to provide a more theoretical angle for my approach in general, taking a broader sample of video game interfaces and generating tools and vocabulary for future work in this direction. Together, this thesis provides the first thorough and multifaceted design of techniques, evaluation, and analysis of the overarching strategy of learning from video games for teleoperation interaction techniques.

For the remainder of this chapter, I provide an overview of each component and its methodology. I discuss the overall approach to how I choose and implement each video game-inspired teleoperation design (or my framework), as well as provide some details and results of my evaluation of those designs. I finish the section with a discussion of the significance of this thesis to the broader community, as well as a summary of this introduction.

1.2 Designing Video Game Inspired Teleoperation Interfaces

Video game interface design includes a wide range of design elements such as interface widgets, layout, menus, and other techniques for simplifying or improving interaction. Video games generally consider player psychology as well; designing for player psychology can encourage changes in behaviour, add motivation, manage stress, or augment storytelling, such as with flow theory (Johnson and Wiles 2003) or gamification (Deterding 2012). The use of such techniques to change player psychology has been noted as a primary difference between everyday software interfaces and games (Langer, Hancock, and Scott 2015).

Taking inspiration from video game interaction techniques, I design, prototype, and evaluate new teleoperation interfaces following standard human-computer interaction methodologies (Beyer and Holtzblatt 1997). In particular, I leverage knowledge that is highly related to common teleoperation problems, such as how to direct attention (Chapter Three, published in Rea et al. 2017), the effects of perception (Chapter Four, published in Rea and Young 2018; Rea and Young 2019a), and sensor visualizations (Chapter Five, published in Rea and Young 2019b). I have found that the effects of such interfaces, such as changes in perception of the controlled agent or mental workload of the operator, can be successfully replicated in my teleoperation interfaces.

My evaluation of these techniques is based on user experiments that have people use my designs and provide feedback. My analysis is human-centered and focuses on user experience, including performance, usability, and psychological measures such as workload, emotions, and operator perceptions. More traditional performance measures include task completion time, collisions, or general awareness (Chen, Barnes, and Harper-Sciarini 2011; Steinfeld et al. 2006), while other user-centered measures focus on ease of use like perceived workload (Steinfeld et al. 2006), and open-ended feedback (J. E. Young et al. 2010). Further, to measure other potential psychological effects seen in video games, such as changing how fast a person believes a vehicle drives [Mario Kart 7, Nintendo, 2011; Hi-Octane, Bullfrog, 1995], I measure how my designs affect the operator's perceptions of a robot's capabilities, such as how fast or safe they believe it is. Such an exploration helps me to understand the nuances of applying video-game inspired interfaces (that were designed to leverage player psychology) to designing teleoperation interfaces. It further allows me to reflect on the feasibility of my approach, while evaluating the practical benefits of the interfaces we explore.

In this thesis, I explore the use of video game inspired teleoperation interfaces and how they can improve teleoperation from several angles, such as directing operator attention or influencing their expectations of robot performance. While specific approaches vary, each project uses a human-centered design process of designing, implementing, and performing human-centered evaluations. I provide an overview of my specific approaches, designs, and results for each project below.

1.2.1 Human Perception and Attention

I designed interfaces to support operator attention management and awareness – interfaces that improved an operator's ability to direct their attention to potentially important areas in a robot's camera feed. As operators are often visually searching camera feeds during operation, such as for environmental dangers in a search and rescue mission, a computational aid to indicate areas to inspect could improve operator performance and reduce the cognitive cost of visual search. These computer algorithms may detect events missed by the operator, and reduce the effort needed to search for them visually. However, an interface design may be too attention-grabbing, distracting an operator during an important task or after they have already evaluated the area of interest. Thus, how to convey such information without interfering with the operator's continued visual search is an open problem that I tackle in this project.

Video game players are often also visually searching a rendered scene, either for items to collect, incoming enemies, their next navigational objective, and more. Game designers have experimented with methods to indicate important things, like a computer-controlled companion that highlights enemies or important parts of the environment [e.g., The Legend of Zelda: Ocarina of Time, Nintendo, 1998; Mass Effect 2, BioWare, 2010], automatically increasing the saliency of interesting areas [e.g., Super Mario Galaxy, Nintendo, 2007; Grand Theft Auto IV, Rockstar Games, 2008], or employ methods to bring the player's attention to an incoming event [e.g., Goldeneye 007, Rare, 1997; Bioshock, 2K Games, 2007; Doom, id Software, 1993]. I take inspiration from common video game interface elements to design attention-grabbing interfaces for teleoperation, and ground my designs in related work on visual attention in psychology, and explore and test numerous different visual attention.

I took an iterative design approach, employing several design, evaluation, and update cycles that integrated the results of my analyses at each step. My evaluation of my attention guides focused not only on operator detection of important events, but also how the interface design may have distracted the operator from their continuing visual search. My video game inspired interfaces were effective at drawing attention, but some designs performed better, and some were considered distracting, increasing cognitive load. I outline the trade-offs of these design techniques in terms of design parameters found in the visual perception literature, discussing the parameters' effects on drawing attention accurately, drawing attention quickly, and being distracting. Further, I provide a set of effective attention-drawing interfaces for use in teleoperation (e.g., Figure 1.2).

This project demonstrated how attention-grabbing techniques from video games can be adapted to teleoperation tasks to balance between being distracting and helping operators search video feeds from robots. On a larger scale, my approach serves as evidence for how learning from video game techniques help inform new interfaces that convey ongoing events to the operator, improve operator performance, and reduce cognitive load.

1.2.2 Priming Operator Perception of their Robot

I aimed to use priming, how I present a robot and its capabilities, to alter people's perception of a robot's physical characteristics to impact driving behaviour. I use the term priming to mean the use of stimuli that remind people of previous behaviours and experiences that then affect their current thoughts and actions (Bargh, Chen, and Burrows 1996). Specifically, I explore whether I can influence an operator's perception of the risk of driving a robot (e.g., if it is faster, more powerful, or less stable), by priming the operator on expectations about what a robot is physically capable of, and investigate whether this may encourage a change in the

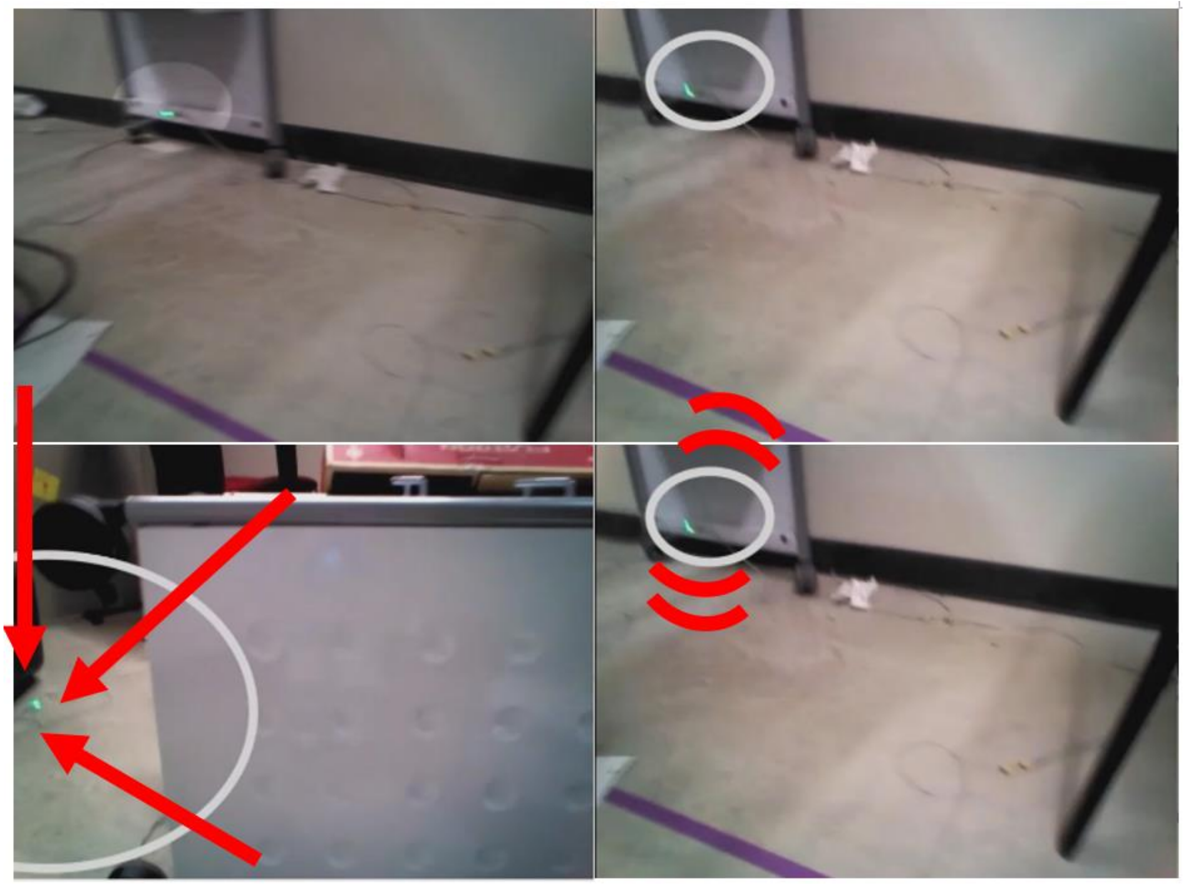

Figure 1.2. A sample of some of the attention-drawing interfaces we tested. From the top-left, these included a darkening effect, a circle, a circle that quickly shrank towards the target, and a circle that bounced around the target.

operator's driving behaviour. For example, I may prime a person to believe that a robot is light and fast, which may remind them of a sporty vehicle. This, in turn, may impact how operators drive the robot (e.g., cautiously as it may be more dangerous).

Video games themselves are designed to engineer and shape player perceptions of entities in-game. For example, video games manipulate the display of player health to make it seem like the player is in more danger than they really are, making the player feel more like a hero by thinking they escaped a dangerous situation [Assassin's Creed, Ubisoft, 2007]. Racing games have manipulated how players perceive a vehicle's driving ability by misrepresenting the vehicle, such as displaying how a vehicle can accelerate faster than another when that is actually untrue [e.g., Hi-Octane, Bullfrog Productions, 1995; Mario Kart 7, Nintendo, 2011, see "Mario Kart 7 In-Game Statistics" 2019]. Video games even change enemy artificial intelligence to make the player feel more powerful, such as by programming enemies to not attack players when the player is looking elsewhere [e.g., DmC: Devil May Cry, Ninja Theory, 2013; Spec-Ops: The Line, Yager Development, 2012]. One goal in all of these examples is to create certain player experiences and behaviours, and they do this by leveraging player psychology to believe things that are not actually true (see more details in Scheurle 2017).

Thus, the goal in this chapter was to explore how I may similarly encourage certain perceptions of a robot or its abilities in order to improve an operator's performance and experience. Specifically, I investigated if I can prime operators to drive more safely by encouraging them to believe their robots were, for example, faster than another, without having to alter the actual robot. I designed, implemented, and tested several techniques for priming an operator about a robot's capabilities. I investigated techniques such as how I described the robot to the operator (e.g., as fast, heavy, etc.), or how the robot felt as it drove (e.g., heavy or slow) based on the joystick and acceleration properties (Figure 1.3). I found that priming can indeed make operators believe a robot is more durable, safe, or even faster than others. In some cases, priming even improved driving behaviour by reducing collisions. My results demonstrate that designers should not only focus on the functional and performative aspects of their robot and interface designs, but also consider how a robotic system is presented to their operators and how that can influence the operator's behaviour and perception of the robots themselves.

Figure 1.3. One priming technique we employed was to provide information sheets to the operator that described different characteristics of the robots they drove, though the information did not reflect actual robot ability.

My investigation into priming demonstrated how video game techniques that try to encourage players to believe certain qualities of the virtual world may be adapted to teleoperation to change the expectations and beliefs operators may have about the robots they control. This provides a powerful way to shape teleoperation perception and behaviour without having to change the robot in any way. It further suggests that my design choices in my interfaces in general may be influencing the operator in unplanned-for ways; understanding priming gives designers a new tool to both change the operator experience and understand why operators may believe or act in certain ways.

1.2.3 Social Techniques for Changing Operator Experience

Video games use social interaction techniques to shape player experience and behaviour. For example, they can be used to build rapport between the non-player characters and the player [e.g. Mass Effect 2, BioWare, 2010], or characters in the game world can talk and bring attention to important ideas and objects in the world [e.g., The Legend of Zelda: Ocarina of Time, Nintendo, 1998], or to communicate performance to the player by having characters complain or appear hurt when the player performs poorly [e.g., Doom, id Software, 1993]. I take inspiration from these uses of social interaction to design a virtual "passenger" that accompanies a robot operator virtually in the interface and reacts to how safely the robot is being driven, like a passenger in a car. (Figure 1.4).

More broadly, it has become common to explore how robots can use human- or animal-like social communication techniques when working with people, in an attempt to improve and simplify communication with them (J. E. Young et al. 2009). For example, autonomous robots co-located with people can use techniques such as expressive movement (Sharma et al. 2013a), gaze (Breazeal et al. 2005), or even animal-like tail movements to convey robotic state or intention (Ashish Singh and Young 2013). However, apart from telepresence (where a robot is a proxy for a person in human-human interaction), there has been little work done that explores how a *teleoperated robot* can similarly use social techniques to support their remote operators. As such, I bridge teleoperation and social human-robot interaction by exploring how a teleoperated robot can use social techniques, adapted from video games, to impact the operator.

I designed, implemented, and evaluated a simplistic virtual passenger that resides in the

a)

b)

Figure 1.4. a) An on-screen "virtual passenger" agent reacts to poor driving by exhibiting anxiety, with the intention of impacting the teleoperation experience.
b) the interface displayed during robot teleoperation.

teleoperation interface. The passenger reacts to the operator's driving using social feedback, for example, by acting scared, in an attempt to make the operator drive more carefully. In a formal user study, I compared two different styles of agent reactions to explore how different social signals may impact an operator during unsafe operation. I found that, depending on the social reactions shown during operation, I could change the operator's experience by influencing their emotions and their impression of their own performance.

By combining video game ideas of how to leverage social behaviours and social human-robot interaction strategies into teleoperation interfaces, I influenced the operator's user experience.

This is particularly important as it may further affect their willingness to use the technology again, or even change their behaviour during teleoperation (Precht, Keinath, and Krems 2017). This initial study is the first to bridge these fields, and only scratches the surface of potential social interfaces; both social human-robot interaction and video games successfully change people's behaviours and experience with different or more complex social interactions, which may also be applicable in teleoperation interfaces.

1.2.4 A Framework of Video Game Interaction

In this thesis I have designed and experimentally verified several video game inspired interfaces and how they can improve or change a teleoperator's behaviour or experience. I created a framework to move beyond individual specific instances and describe a broader sample of video game interaction to more generally comment on the applicability of my approach of improving teleoperation by learning from video game design. To this end, and to facilitate continued work in the area, I aimed to provide the community with a simpler, abstract representation of the techniques used in video games and the problems they aim to solve.

To construct my framework, I conducted a survey of different video game interaction techniques. I analyzed common techniques and goals in video game interaction design, resulting in a description of video game techniques from four angles: 1) what video game designers intend to communicate to players, 2) what interface techniques have been designed to communicate those things in-game, 3) general implementation strategies, and 4) high-level user experience goals that guide interaction design. This enabled us reason at a more abstract level about the applicability of broader trends in video game interaction, as well as map existing game interaction techniques for future work.

There are thousands of video games, each with a range of interfaces to potentially explore in teleoperation. Because of this, I chose my methods and data to provide a sample of the interaction space; I sampled from game ranking lists, playing games when available, or observing play (e.g., online videos), and observing interaction techniques used in those games. My survey sample targeted games that share similarities with teleoperation, specifically controlling a remote entity and completing objectives in a virtual environment.

My observations resulted in a rich and complex qualitative dataset; I analyzed the data with the goal of simplifying and distilling seemingly disparate interfaces into groups that share goals or techniques. Specifically, I used iterative open coding (Berg 1989) and axial coding (A. Strauss and Corbin 1990; A. L. Strauss 1987), resulting in groups of high level themes I structured into my framework. My framework simplifies a portion of the large world of video game interfaces into a more tractable abstraction, and it can be leveraged by teleoperation researchers to learn from and explore a new range of video game inspired teleoperation interfaces.

To provide further validation of my broader approach, I considered each result describing video game interaction (e.g., design goal, implementation strategy) and how it can be applied to teleoperation. I further linked notable examples of human-robot interaction research that share similarities to video game interaction design when possible, providing additional experimental evidence for my approach. The broad potential applicability of game-like techniques in my framework demonstrates that formally looking to interaction techniques from video games can be applied to teleoperation more generally.

1.3 Significance

Teleoperation has many potential applications in both the home, industry, and other areas such as search and rescue. However, safe and efficient teleoperation is not trivial for operators due to the complexities of controlling robots (including network latency and robot designs with many sensors and joints) and understanding the remote environments with limited or obtuse sensor data. I investigate how the human-robot interaction community can improve teleoperation interfaces in practical scenarios by learning from video game interfaces, taking four different angles: directing operator attention, priming perceptions of robot capability, using social agents to influence operator experience, and creating a framework of video game interaction. The framework classifies a breadth of video game interaction techniques, provides tools and vocabulary to discuss and apply those techniques to teleoperation, and can help highlight future work for creating new teleoperation interfaces.

Much of my contribution is in my exploration of different high-level video game interaction design approaches and how I adapted them to teleoperation interfaces and experimentally evaluated their effects. In some cases, I contributed whole new approaches to teleoperation interface design, such as through influencing operator perceptions with priming, and integrating social techniques in video games and social human-robot interaction into teleoperation.

My research demonstrates that video game interaction designs can improve the operator experience, and that operator experience itself can improve task performance. I further showcase how these game techniques can benefit teleoperation by leveraging knowledge of human psychology for a range of effects. Future improvement to teleoperation interaction design can use video games as a springboard to influence operator experience and focus on how those designs leverage natural human mental processes such as perception, emotions, and social interaction.

The results from my game survey and analysis provides the HRI community with a bridge from the video game industry to teleoperation. It provides tools and vocabulary to help researchers discuss and apply different video game interface designs to teleoperation, as well as a set of novel interaction techniques and directions for future exploration. By considering my own video game-inspired designs and my survey together, I provide the first thorough and multifaceted analysis of the overarching strategy of learning from video games for teleoperation interface techniques.

1.3.1 Contributions

I made several contributions to the teleoperation interaction design field in human-computer interaction and human-robot interaction:

1) I designed, prototyped, and evaluated a set of concrete interaction techniques that can be immediately applied to teleoperation. This includes a set of successful attention-drawing designs, three methods for priming perceptions of operator safety, and a method for improving operator's emotions during teleoperation.
2) I successfully demonstrated that video game inspired interfaces can influence operator behaviour and perception in teleoperation interfaces. This includes verbally and visually describing the robot's abilities, suggesting robot abilities through a tangible

feel, or modifying the robot's driving profile to appear safer and easier to drive.

3) I explored how social interface designs inspired from video games and social robotics can be leveraged to shape teleoperation experience. This pioneers a bridging of social robotics, teleoperation, and video game fields, and opens teleoperation interface design to take advantage of social techniques from each field.

4) A framework that describes video game interaction from several angles for the purposes of discussion and application to teleoperation. This enabled us to establish the similarity between the problem spaces and design goals in teleoperation and video game interaction design. It further provides ways to simplify and discuss the wide variety of game interaction designs for future application to teleoperation.

5) I provide a set of design parameters for directing visual attention that can be used in other visual interface designs in teleoperation.

6) A reflection on how my video game-inspired designs' focus on improving user experience can improve operator performance, and how user experience itself is intrinsically valuable to the field of teleoperation.

In the remainder of the thesis, I explore the background and foundational work in the area, including a survey of relevant related work (Chapter Two). Following, Chapters Three to Six detail the specific projects as indicated above. I finish with a high-level discussion of the thesis as a whole (Chapter Seven), and the limitations to my work (Chapter Eight), followed by contributions and concluding statements.

Chapter Two: Interface Design, Teleoperation, and Video Games

In this section I give a thorough background of human-robot interaction, highlighting how the work in this thesis fits within the broader landscape of robot interface work. As part of this, I highlight how human-robot interaction fits under the umbrella of human-computer interaction and interaction design in general. In particular, I narrow in to consider how human psychology is intertwined with interaction design, and how other work has looked at the overlap between gaming and interface design.

2.1 Interaction Design

Interaction design is increasingly thought to be an important and fundamental aspect of all parts of our life (Preece, Sharp, and Rogers 2015). The physical shape of an object can suggest how to use it, like the handle on a mug or door (Gibson 2015; Norman 2013). The design of roads can change traffic flow and safety for those who drive on them (Elvik 2007). Programming language design can help express certain ideas or solve certain problems more quickly and easily (Kato, Sakamoto, and Igarashi 2013; Kato and Igarashi 2014). Thinking of how people perceive, think about, and use things in the world can help design tools that are more pleasant to use, inspire better communication and understanding, finish tasks faster, reduce injuries and mistakes, and more (Preece, Sharp, and Rogers 2015). Interaction design is large and far reaching, but even within technology design alone, it encompasses aspects of human factors, ergonomics, information processing, and social and cognitive psychology (Preece, Sharp, and Rogers 2015). The following sections will highlight different components of interaction design and how they can affect our lives.

2.1.1 Human Factors and Ergonomics: the usability of the physical world

Human factors and ergonomics are a part of usability research that focuses not just on how people use technology, but how the people, their environment, and the technology all affect each other (Murrell 1965). While emphasizing physical human needs and limitations, mental processes such as workflows are also an important consideration that could impact performance (Murrell 1965). This further expands to include parts of psychology (Grandjean 1980), and generally focuses on psychophysiology (how mental state can affect things like heart rate, blood pressure, etc.), behavioral-cognitive models (e.g., such as workload, awareness of the environment and task), and teamwork related sociology (Stanton et al. 2006). One core idea is that, by considering the needs of a human body and mind, we can design better workplaces, tools, and processes to improve efficiency, comfort, and safety (Stanton et al. 2006; Sharit 2006; Murrell 1965).

Some strategies for improving design have emerged from this field generalize to technology as well. These strategies focus on how fundamental aspects of our physiology, psychology, or environment affect us in a way that is reflexive or otherwise difficult to avoid. One important approach is to understand the physical limits of people and where errors may occur (Sharit 2006). For example, people may only be able to visually keep track of a certain number of objects at once (Thornton et al. 2014) which could limit their ability to react to unsafe events quickly. We can also anticipate possible problems by understanding processes, or how people process information, such as understanding how the height or field of view of a driver can affect how fast they think they are driving (Chen, Haas, and Barnes 2007). Instead of correcting for potential weaknesses, designers can also focus on leveraging natural strengths of human abilities, for example leveraging how human peripheral vision is good at detecting motion (e.g., to improve visual attention, Daniel J Rea, Seo, et al. 2017, or spatial awareness, Seo, Young, and Irani 2017; Gustafson et al. 2008). Thus, it is important to understand these human tendencies on both a theoretical level – such as the limits of reaction time – and practical level

– how the theory can be leveraged to understand and design tools and workflows in the real world.

2.1.2 Cognitive Processes and Workload

While the study of mental processes is part of Human Factors, understanding how people think consciously and rationally is itself a broad and deep subfield that has developed its own theories and methods. Cognitive psychology is the study of how people perceive, process, store, and recall information (Neisser 2014). For my purposes in interaction design, cognitive processes are important as they can both inform how interfaces should be made (adapting interfaces to people) and help us understand how they are used (how people adapt to interfaces).

Cognitive processes manifest themselves in interaction design in many ways. For example, we can design better layouts for web pages by understanding how our brain uses visual properties to infer relationships between objects and visual information (e.g., objects close together are assumed to be related according to gestalt principles; Koffka 2013). The famous "seven plus or minus two" rule (Miller 1956) describes how much new information a person can concentrate on at once, and can help guide designs of things used in learning (e.g. textbooks or presentations), attention (tracking multiple objects), or short-term memory. Software is often concerned with aiding people with understanding or working with knowledge, and thus cognitive psychology plays a part in software design (Preece, Sharp, and Rogers 2015). As teleoperation also needs operators to perform complex mental tasks, such as building an understanding of the surrounding area (Chen, Haas, and Barnes 2007) or reasoning about how to best control a robot arm (Leeper et al. 2012), understanding and aiding cognitive processes is important to teleoperation as well.

Mental Workload

One concept in cognitive design that is particularly relevant to our work with robot teleoperation is known as *workload*. *Mental* workload is "the relation between the function relating the mental resources demanded by a task and those resources available to be supplied by the human operator" (Parasuraman, Sheridan, and Wickens 2008, p. 145; Wickens 2002). Thus, by definition, it is related to how and how much information a person is perceiving, processing, and reacting to, and is one aspect of cognitive psychology. The definition of workload sometimes contains a measure of physical stress as well due to the mental resources required to handle physical coordination (Sheridan and Simpson 1979).

Workload itself is a part of usability and not performance; for example if two people are doing the same task at the same performance level, one may be handling the task easily without tiring while the other is concentrating very hard (Parasuraman, Sheridan, and Wickens 2008). If a person cannot supply the mental resources required to keep up with mental demand, however, it can cause decreases in performance (Parasuraman, Sheridan, and Wickens 2008; Wickens 2002; Bailey and Iqbal 2008; Sheridan and Simpson 1979; Burke et al. 2005). Also, constant high demands on mental resources (e.g., continuous high workload) is known to become more difficult over time, and thus, like physical work, it cannot be assumed that a person can continually maintain a level of performance indefinitely, even if it is initially handled with ease (Sheridan and Simpson 1979; Parasuraman, Sheridan, and Wickens 2008).

The importance of workload can be seen in many areas. It is known to be related to safety and performance in the study of traffic psychology (Fuller 2005), aircraft piloting (Sheridan and Simpson 1979), and control of other machines such as robots (M. a. Goodrich, Crandall, and Barakova 2013; Chen, Barnes, and Harper-Sciarini 2011). Thus, measuring workload becomes important. Physiological measurements are possible, such as galvanic skin response (Stanton et al. 2006; Sharit 2006), brain activity (Lim et al. 2010), and even eye movement (Van Orden et al. 2006). However, many of these physical responses may be confounded due to task requirements (e.g. a task could require an easy visual search that increases eye activity but not greatly impacting workload), or have potentially overlapping causes (changes in galvanic skin response occur for many reasons; Montagu and Coles 1966). Thus, experiment participants filling out questionnaires after a task is a common method (e.g. Hart and Staveland 1988; Sheridan and Simpson 1979). These have the benefits of being easy to administer, but have many drawbacks such as the limitations to drawing conclusions about causal relationships due to being unable to verify if questionnaire responses are true, or distorted due being a subjective interpretation of a feeling by the person answering the questions (Razavi 2001). Nevertheless, workload questionnaires have proven to be a useful source of information if these limitations are taken into account (Hart, Sandra 2006), and have also been proven useful in understanding teleoperation interface design (e.g., Daniel J Rea, Hanzaki, et al. 2017; Chen, Barnes, and Harper-Sciarini 2011), and video game interface design (Koeffel et al. 2010).

2.1.3 Social and Emotion Psychology for Interaction Design

Modern interaction design research increasingly considers other parts of human psychology when designing and evaluating interfaces. Research has found that studying how use of technology affects a user's emotions (or *affect*) or social relationships can impact both use and perceptions of technology (Norman 2004), including video games (e.g., Jennett et al. 2008) and robots (J. E. Young et al. 2009). I outline below how affective and social components can affect and aid interaction design.

User Affect

Technology can affect a user's emotion, or take a user's emotions into account to make decisions (C. Peter and Beale 2008; Christian Peter and Herbon 2006). Indeed, even from a purely performance-focused perspective, research into emotion has found it can improve or reduce effectiveness at cognitive tasks (B. Fredrickson 1998; Isen 1987). Thus, changes to a user's emotional state may affect teleoperation, which itself is a mentally demanding task (Steinfeld et al. 2006).

Aiming to change a user's affect can be the goal of a technology on its own; overall positive emotional state can be attributed many social, physical, and mental benefits (B. L. Fredrickson 2001). For example, positive emotions can make people consider more actions to solve problems (B. L. Fredrickson 2001), or reduce risk of coronary heart disease (Blascovich and Katkin 1993). Negative emotions, conversely, may negatively impact the ability to think logically and clearly (Gross 2002) which can lead to major safety issues (e.g., while operating machinery; Precht, Keinath, and Krems 2017).

Targeting certain emotional responses when using technology has seen other benefits as well, such as changing the interpretation of data visualizations (Harrison et al. 2013), or resulting in increased interest in and engagement with software learning (Langer, Hancock, and Scott 2015). In human-robot interaction, detecting or influencing human emotion (e.g., Erden 2013; Riek, Paul, and Robinson 2010) or expressing emotion (e.g., Feldmaier, Stimpfl, and Diepold

2017; Sharma et al. 2013) using different social techniques is considered a fundamental part of the interaction (J. E. Young et al. 2010; J. E. Young et al. 2009). Thus, understanding how technology influences or interacts with people's emotions is important for developing interfaces that keep users in a positive mindset and engaged with their task.

Social Psychology

Humans are social creatures, and we interact with technology socially as well (Lee and Nass 2010). Technology can change how we interact with other people, such as how the height or internet connection of a telepresence system can change how people local to the robot perceive the remote human operator (e.g., Kristoffersson, Coradeschi, and Loutfi 2013). Further, the design of forums in online classrooms can affect how and how often students interact with each other (Wang et al. 2015). Autonomous robots can leverage knowledge of social structures to affect the relationship between two people (e.g., D. Sakamoto and Ono 2006). Thus, by considering technology in the context of people and their relationships with each other, that same technology can be designed to improve human-human interaction.

When we take into account that people interact with technology inherently socially, human social abilities and social tendencies can be incorporated into technology design to improve interaction experiences (Reeves and Nass 1996; Lee and Nass 2010; Breazeal 2003). For example, technology may even mimic social skills to help itself in an interaction (e.g., with software agents: Forlizzi et al. 2007; or robots: Young et al. 2009). Even seemingly simple understanding of body language like gaze and pointing can improve human-robot interaction (Breazeal et al. 2005). The social nature of the human experience must be considered in interaction design (Reeves and Nass 1996; Lee and Nass 2010; Preece, Sharp, and Rogers 2015), and, as seen in the above examples, may even serve as the basis or inspiration for new ways for people to interact with technology.

2.1.4 Interaction Design in Human-Computer Interaction

Human-computer interaction is the study of interaction design with regards to people and computing technologies. (Dix et al. 2003; Preece, Sharp, and Rogers 2015). As the high-level goals of human-computer interaction are shared with interaction design (Preece, Sharp, and Rogers 2015), the two fields share many of the same approaches and evaluation methods.

Computer and software design can benefit from human factors, cognitive evaluation, affective design, and social techniques – components of interaction design outlined above. For example, software can apply strategies from human factors to increase task performance (e.g., Pascoe, Ryan, and Morse 2000), or understanding human cognitive processes and workflows can help improve software-aided prototyping and design (e.g., Mine, Yoganandan, and Coffey 2014; D.J. Rea, Igarashi, and Young 2014). Interactions that take into account or influence user emotion has spawned the entire subfield of affective computing (Picard 2000, e.g., Sharma et al. 2013; Forlizzi et al. 2007; Jennett et al. 2008). Social psychology has also found to be applicable to technology in the *computers-as-social-actors* (CASA) paradigm, which stipulates that people tend to treat technology as if it could respond to social interaction (Reeves and Nass 1996).

Thus, considering how both the human body and mind works is important for improving usability, performance, and comfort of software and computing hardware. Below, I detail some concepts from human-computer interaction that are particularly relevant to this thesis.

Psychology and Software

As my approaches in this thesis focus on leveraging psychological theories and phenomena, I wish to further emphasize how human-computer interaction has long leveraged psychology research to inform its research and designs. One example is how research on visual perception gave forth the idea of *affordances*, or aspects of design that suggest their use to an observer (Gibson 2015, e.g., a doorknob looks like it can be turned), which has been carried over and extended into software design to improve usability (Ware 2012).

Affordances in software are more nuanced and are an evolving concept as many actions are abstract and virtual (Mcgrenere and Ho 2000). As an example, in Gibson's (2015) and Norman's (2013) classical interpretation of affordances, hyperlinked text does not physically afford the action of clicking to bring you to a new page; that you can click blue and underlined text is a learned convention. However, as the concept spread through the interface design community, such *perceived affordances* can be argued to have become the default interpretation of a software affordance: to be perceived as a method for a user to take an action, such as how some virtual buttons look *clickable* (Mcgrenere and Ho 2000; Gaver 1991). I note here that psychology, or in this case cognitive processing, is a core theme to how these concepts, and the ones below, are relevant to software.

Other examples of considering psychological theories in HCI include how mental models of tasks or systems can affect interaction (e.g., Forlizzi et al. 2007; Labonte, Boissy, and Michaud 2010; Skalski et al. 2011; Bowman, Koller, and Hodges 1997; Paymans, Lindenberg, and Neerincx 2004), or how visual representation can change how users try to engage with software (e.g., Ullmer and Ishii 2000; Norman 2013; Maier and Fadel 2009; Matthews 2007). Game developers and researchers have also focused on considering the user's psychology by expanding on ways to improve user engagement such as gamification (Hamari, Koivisto, and Sarsa 2014; Li, Grossman, and Fitzmaurice 2012; Deterding 2012), or enjoyment of software by leveraging flow theory from psychology (Johnson and Wiles 2003). I discuss more focused applications of psychological theories to video game and robots below; I emphasize our work continues this approach by leveraging related knowledge from psychology to understand and design new interfaces.

2.1.5 Embodiment

To be part of interaction, people, computers, and other phenomena need to take on an embodiment. This is the physical properties something has, and defines the way it exists and can participate in the world: speech is embodied by sound waves, and any person or system that wishes to interact with or by speech must take this into account (Dourish 2001). This concept has broad implications, revealing that physical forms dictate how we perceive, think about, and learn (Iacoboni 2009; Klemmer, Hartmann, and Takayama 2006). For example, a hammer needs a handhold to be picked up and swung, and a flat hard surface to hit a nail. The hammer designer should also consider the embodiment of the people who will use it: the hammer's handhold should consider the shape, size, and movement constraints of a human hand that will hold it. A person will also think with their own embodiment: they perceive the handle of the hammer based on their own hands and what they know is easy to pick up, and feel the weight and strength of the hammering end with their sense of touch or muscles. By using the hammer, or by watching someone swing a hammer, a person better understands how it can be used – we can think (Klemmer, Hartmann, and Takayama 2006; Sudnow 1993) and can learn (Iacoboni 2009) with our embodiment. Thus, the embodiment of something needs to be considered in interaction design.

This concept applies to human-computer interaction as well (Klemmer, Hartmann, and Takayama 2006). A keyboard's layout takes into account the shape and size of a human hand, as well as the number and movement range of the fingers. An interface should consider people have two eyes that focus on one place at a time, implying limits to multitasking. Thus, a person's physicality should be considered when designing interfaces.

Embodiment extends into our emotional or social interactions as well. What we are thinking is embodied in our vocal tone or our posture; we communicate with others by using our bodies: waving, pointing, miming, etc. This informs how we interact with both computers (Lee and Nass 2010) and robots (J. E. Young 2010). For example, if a robot extends a single gripper or an on-screen graphic is shaped like an arrow, a person will likely interpret this as pointing – the person interprets the gesture by equating one they would make with their own embodiment. I consider how video game designs leverage embodiments, including both the person playing, and how the games virtually embody actions and characters in their virtual worlds, and explore how to apply those ideas to teleoperation interfaces.

2.1.6 User Experience

Looking at the previous sections, we can see that 'usability' is a core theme – how easy is something to reason about, hold, or perform tasks with. However, ideas like user affect and social relationships and influences in technology are more difficult to incorporate into the idea of usability design. This brings us to the term of user experience, which has been noted to be vague in its usage in both the press and academia (Hassenzahl and Tractinsky 2006; Forlizzi and Battarbee 2004). As a beginning point, I quote Forlizzi and Battarbee:

> What is unique to design research relative to understanding experience is that it is focused on the interactions between people and products, and the experience that results. This includes all aspects of experiencing a product — physical, sensual, cognitive, emotional, and aesthetic. (Forlizzi and Battarbee 2004)

Thus, the idea of user experience is necessarily large and nuanced. For example, research has explored the differences and interconnections between ease and fun, ultimately concluding how interaction design should consider both (Carrol and Thoma 1988). Modals of user experience discuss how interaction can be a layered experience, working on mechanical, contextual, and visceral (emotional, instinctual, or biological) layers (J. E. Young et al. 2010). Experience is both linked to the context of an interaction, and the time of the interaction (Hassenzahl and Tractinsky 2006) – if a person's mood or the weather changes, even later in a day, their interaction with something such as robot or word processing software could change. While it is effectively impossible to take into account all aspects of the time and context of an interaction, the concept of user experience encourages us to embrace not just what is functional or usable in our designs and evaluations, but also what is emotional, intuitive, and complex in its relation to humans.

The idea of experience is very important to video games as well (Nah et al. 2014; Hochleitner et al. 2015; Koeffel et al. 2010). Fun is one experiential factor (Carrol and Thoma 1988) and is often a goal of games (Skalski et al. 2011). Games can try to evoke a feeling of climax during a confrontation with evocative music [Overwatch, Blizzard, 2016], or try to make each piece of equipment *feel* unique [Destiny 2, Bungie, 2017]. Games also use a desired experience as a design goal that can impact gameplay design or aesthetics heavily (Sakamoto, 2015). If a game creates a cohesive and effective experience, it can help players become immersed in the game (Ermi and Mäyrä 2005; Mekler et al. 2014). Due to the importance of the concept of experience in both games and interaction design as a whole, I consider it during my designs, and attempt to measure non-usability or non-functional aspects of my interaction designs as part of my evaluations.

2.1.7 Summary of Interaction Design

In the past sections, I have touched on different components and approaches to interaction design in different levels of detail, and I have conveyed the complexities possible to consider in design. Not only should we consider physical aspects such as how comfortable and easy it is to engage in an interaction, we should consider both how to leverage, and the potential effects of, human emotions, social tendencies, and cognitive abilities of people in an interaction. Thus, interaction design is a deep and multi-faceted field that considers both physical and mental components of the human condition, and this background knowledge will ground my approach as I extend teleoperation interaction design by learning from video game interaction techniques.

2.2 Human-Robot Interaction

The field of human-robot interaction studies robotic systems and how they may be used or affect people who use them and are around them (Michael A Goodrich and Schultz 2008). On an abstract level, the field is composed of the two roughly distinct focus areas of social human-robot interaction and teleoperation (J. E. Young 2010). These areas can be defined by the distance the interaction takes place, and how autonomous a robot is. Generally remote robots are controlled by teleoperation (sending commands from another computer such as move forward or pick up an object), or perform routine and unsupervised tasks such as vacuuming. Once robots share the same space as people, it becomes easier for robots to use social techniques. This includes people to give commands to the robot using natural language or body language like pointing and gesturing, or converse with the robot acting as an autonomous social agent. I discuss both but treat teleoperation in more detail due to being somewhat more relevant to this thesis.

2.2.1 Teleoperation

Teleoperation is the act of controlling a robot that is separated from the person controlling the robot (the operator, Michael A Goodrich and Schultz 2008). Problems in teleoperation generally focus on how to control the robot, which may have multiple sensors, driving modes, even complex arms or other instruments. How to enable control of such robots is a complicated interaction design challenge, but in order for a human operator to even make decisions about how to control robots, they must understand what a robot can do and the environment around it; this large problem space is known as situation awareness (Chen, Haas, and Barnes 2007; Endsley 1988). These problems are non-trivial; researchers have been exploring different parts of the problem and designing and improving solutions as part of a long-term research agenda (Endsley 2016).

When designing and evaluating interactions for teleoperation, standard human-computer interaction methods can apply, but robots have specific, unique interaction challenges (Chen, Haas, and Barnes 2007; J. E. Young et al. 2010), necessitating a new field of research. These include staying aware of the robot's state and the environment around the robot (situation awareness), control interfaces that take into account the noisy, slow, and physically difficult nature of robots in the world, and how to provide feedback to the operator. I delve into each of these problems below.

Problems from the Physical World

Due to the physical nature of robots and their tasks, inevitably problems arise due to physical limitations. For example, commands sent to robots are necessarily delayed due to the time needed to send the message to the robot over a network (Figure 2.1), the time needed for the robot to process this command, and then to engage motors and carry out the action. All feedback from the remote area or robot goes through a similar, but reverse, process when being conveyed to the operator. Contrast this to software, which is typically instantaneous from the user's perspective. Control has other problems, such as accuracy (the robot judging if it has succeeded at a command – Susilo 2015) or safety.

Further, traditional software often has very little noise or uncertainty involved. In contrast, robot sensor data has a substantial error component and noise, adding to work for the computer or operator (or both) to process, introducing uncertainty into teleoperation tasks. This difficulty, while not unique to robots, is a problem that may not be physically or financially feasible to correct. Error can be estimated (Smitch and Cheeseman 1986; Susilo 2015), or innovative interface design can try to communicate sensor or task uncertainty to the operator (e.g., Feldmaier, Stimpfl, and Diepold 2017).

Another option is to hide these physical problems from the operator, such as by ignoring or adapting commands when it is unsafe to follow the original command (Cui, Gao, and Guo 2012). Commands may also be simulated by predicting a operator's future commands based on current robot motion to reduce the perceived latency (Zanaty, Brscic, and Frei 2008). These techniques give the illusions of fast or safe control, but if mistakes are made by the robot adjusting or creating its own commands, the illusion disappears, and the system or operator will need to correct the mistakes, potentially frustrating them and causing higher workload.

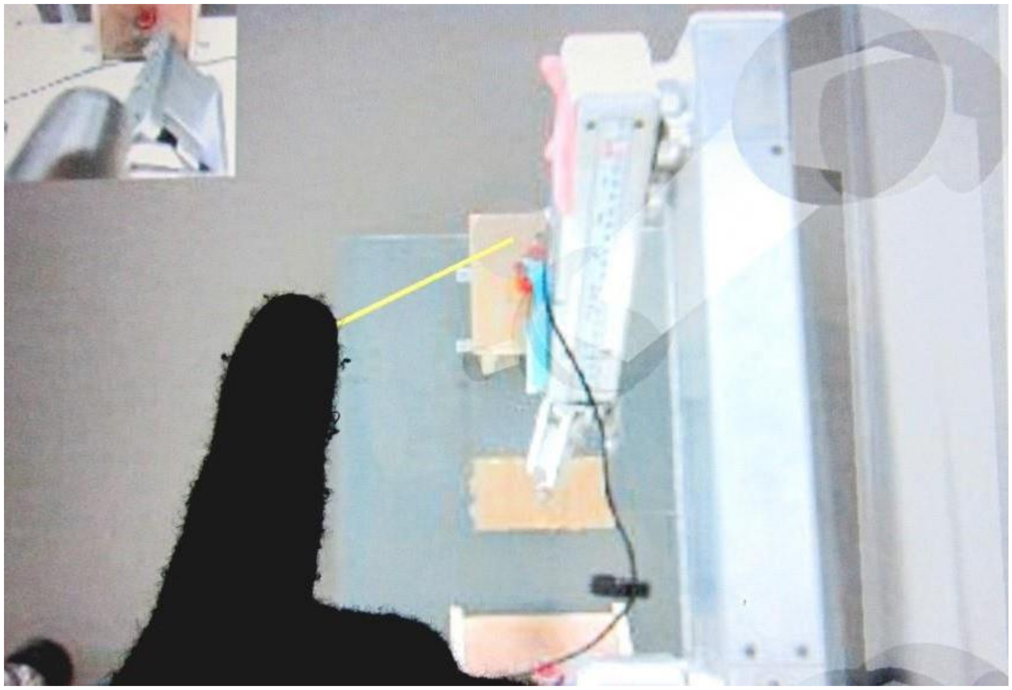

Figure 2.1. A robot arm is commanded to move left. To help the operator understand the result of the command, a virtual model of the robot arm simulates the move faster than the real robot arm moves. From (Ashish Singh et al. 2013).

While both latency and error are common in robotics applications, I found fewer interaction research projects focusing on these issues, with improvements often coming from improved hardware. Video games may be a potential source of solutions for this issue; while video games are generally purely virtual and do not need to worry about many of the problems in the physical world, online games still have network latency. This is especially challenging real-time in multiplayer online games, such as competitive shooting games where players face each other in gun combat where success relies on knowing precise locations of opponents and quick and accurate reactions give an advantage [Overwatch, Blizzard, 2016; Destiny 2, Bungie, 2017, etc.]. These games also use prediction of player movements and clever presentation techniques presenting a real-time façade to users. Thus, video game interfaces may have interaction solutions to some of these less studied problems.

2.2.2 Remote Feedback and Situation Awareness for Teleoperation

In order to control a robot remotely, react to changes around the robot, and interact with people near the robot, detailed information from the remote area must be sent back to the operator. An operator needs to know about their robot – states such as battery level, or robot pose, for example, if an arm is extended. The operator should also understand where the robot is in the remote environment and be aware of where other people and objects are relative to the robot. An operator should also be aware of events that have happened or are happening now, what actions they can take with their surroundings, and what the results of those actions may be. In other words, an operator must be aware of their robot's embodiment (Section 2.1.5).

Considering this amount of information at once creates a cognitive burden for the operator, in addition to any other task-specific actions or goals they may have. Collectively, the problem of staying aware of all these aspects of a situation fall under the umbrella of *situation awareness* (Endsley 1988). Situation awareness can be used to examine problems in many areas, and in teleoperation can be understood as the knowledge that both the remote robot and humans have of the others' intents, commands, location, and activities.

Due to its large scope and relevance to teleoperation (Endsley 2016; Chen, Haas, and Barnes 2007; Yanco and Drury 2004a), I delve into different parts of the situation awareness problem in teleoperation in this section. One main problem unique to the remote-control aspect of teleoperation is that building and maintain awareness of the remote environment can be difficult for an operator due to the lack of their familiar and normal human senses. For example, a person generally understands their own physical pose with proprioceptive senses or senses of motion, which must be communicated by the teleoperation interface in some way. Some senses, like vision, are common to be sent from a robot, but are more limited than a person's abilities, such as how wide an area can be seen (Chen, Haas, and Barnes 2007; Endsley 1988). Similar to the operator's awareness of the robot, the robot should be aware of the operator, but is limited to its programming and whatever commands (and rarely, sensor data) are sent from the control area. The majority of research I found focused on the problem of conveying the remote robot's and environment's state (e.g., Chen, Haas, and Barnes 2007), but it is also important that the remote robot understand the human state (e.g., goals, intent, the meaning of control inputs, etc., Endsley 2016).

Operator Awareness of the Robot and the Remote Area

Teleoperator awareness of the remote environment around a robot typically relies on a number of sensors on the robot (M. a. Goodrich, Crandall, and Barakova 2013). This sensor data must be conveyed to the operator, and in turn be interpreted correctly by the operator in order to understand the state information, predict the future state of the situation, and make good decisions (Endsley 1988; Yanco and Drury 2004b; Endsley 2016, see Figure 2.3). This makes many design decisions task- and robot-dependent; for instance, egocentric (first-person) cameras can provide a better detailed view (Seo et al. 2017), while exocentric cameras (third-person) provide a better sense of where the robot is in its environment (Saakes et al. 2013; Seo et al. 2017). Which sensors are available to an operator and how they are presented to the operator are important design choices for teleoperation.

While the use of normal human senses is limited in the remote environment, robots can be equipped with a number of sensors that can improve upon or imitate human senses. However, due the embodiment of the system, some of this data is difficult to convey with only visuals or sound that computers can produce. For example, robots may have motors that detect forces applied to them. Normally, people feel resistance in their muscles or feel their body being pushed. It is not obvious how this information should be conveyed on a computer (Reveleau et al. 2015, see Figure 2.4). Sensors can be continued to be added to increase functionality and ways to understand the remote world; in addition to force information, awareness can be improved by greater spatial understanding from depth sensors that can see in the dark (Mast et al. 2015), or we can understand proximity of nearby objects with sonar (Nielsen, Goodrich, and Ricks 2007). While these results may make it appear that adding more sensors will improve operator performance, it is not trivial – numerous sensors and data alone are not useful to an operator due to the limited amount of information a person can process in real time (Drury, Scholtz, and Yanco 2003). Thus, the presentation of the sensor data needs to be considered to improve operator performance and reduce workload (Yanco and Drury 2004b; Drury, Scholtz, and Yanco 2003).

This idea of data presentation to reduce workload includes clever visualizations of sensor data to reduce complexity or overlays of calculations on the video feed. For example, the location data from other robots around the teleoperated robot can be presented in a way that intuitively summaries position and distance (Seo, Young, and Irani 2017). Sensor data from multiple sources can be integrated in such a way that it better illustrates relationships in the data. Camera video feeds can be placed within each other when they overlap to better represent how each camera is oriented (Seo et al. 2017), or within a map the robot is generating (Nielsen, Goodrich, and Ricks 2007). Layouts that combine sensor data in such ways are essentially performing situation awareness processing for the operator, instead of having the operator relate data from separate visualizations themselves.

Instead of processing data into a visualization for the operator to interpret rationally, feedback from robots can be presented in a way to take advantage of how people think naturally. For example, collision avoidance feedback can be integrated into the physical controller of the

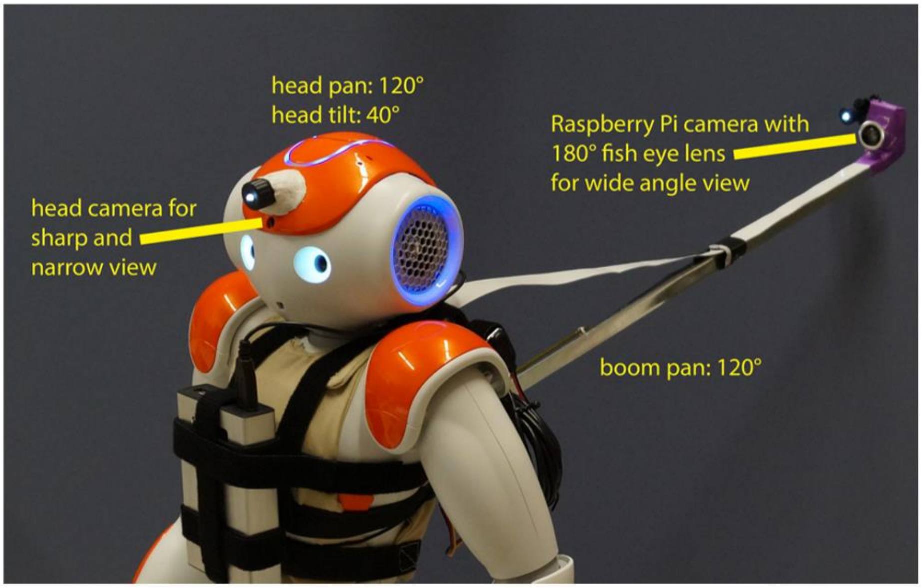

Figure 2.3a) A robot can have multiple cameras and other sensors, and numerous joints and limbs. Conveying all the sensor data and robot joint states to the operator is a difficult problem in situation awareness. Image from Seo et al. 2017.

Figure 2.3b) The view from the first person camera of the robot above. The robot's head can rotate independently from the robot's body, and is turned slightly to the left. If the operator tells the robot to move towards the orange object by pushing forward on a joystick, the robot will actually move in the direction pointed to by the triangle. Understanding robot state and how that affects the robot's controls is another problem in situation awareness. Image from Seo et al. 2017.

robot, such as with a force feedback joystick that pushes the operator's hand in the direction needed to avoid the obstacle (Hacinecipoglu, Konukseven, and Koku 2013); this leverages the force sensing and touch parts of our embodiments. Interfaces can be designed to leverage human cognitive processes, such as displaying information in ways to naturally draw an operator's attention (Teng, Kuo, and Tara 2013). Alternatively, designs can take advantage of how humans process social information, such as how movement can be designed to communicates state, in a sort of robot body language (Feldmaier, Stimpfl, and Diepold 2017; Sharma et al. 2013b). Moving forward, I explore this idea of leveraging natural mental processes, and recommend it as an area with potential impact in teleoperation design (Chapter Three, published in Daniel J Rea, Seo, et al. 2017, Chapter Five, published in Daniel J Rea and Young 2019).

As shown in interaction design, which data a robot sends back and how that is presented to the operator should consider physical embedment, human factors, and mental processes. By taking all of these into account, designers can improve teleoperator performance and experience.

Robot Awareness of the Operator

As a complement to helping operators understand a robot's state and remote environment, situation awareness also includes helping robots be aware of the operator. In other words, the algorithms that make decisions for the robot should, to properly execute commands and display appropriate feedback, consider the commands from the operator as well as how those commands serve the operator's goals (Endsley 1988). For example, a robot may understand an operator wants to go forward, but that could hit an obstacle to the robot's side. The robot can provide subtle feedback to a haptic controller, nudging the operator's hand to the side to encourage them to send a small turn command (Hacinecipoglu, Konukseven, and Koku 2013). Thus, by understanding the operator's goals, a robot can autonomously adapt commands to a dynamic and potentially dangerous world in an intelligent way.

A robot may also be aware of an operator's state, such as by user modelling – algorithmically predicting what a person is thinking, feeling, or will do. For example, control can be simplified by guessing operator intent, as suggested above, and using automation to move towards that goal (e.g., Gopinath, Jain, and Argall 2017). Robots can also consider *how* an operator thinks

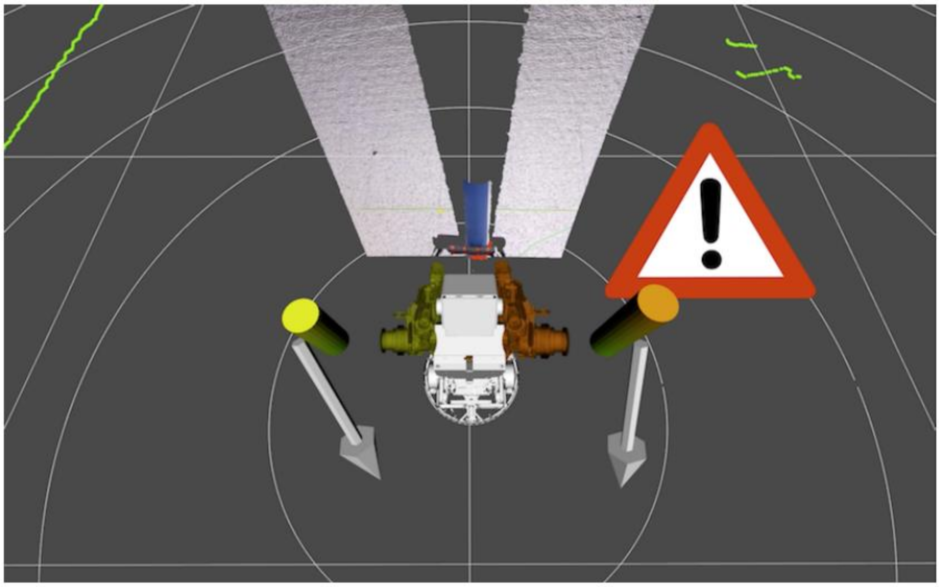

Figure 2.4. Arrows visualize the force being applied to pole objects. People normally feel such forces physically, so it is difficult for humans to intuit forces simply from visuals. Taken from Reveleau et al. 2015.

a task should be completed as input to automation algorithms can improve performance and make control easier, such as asking the operator to give input when an algorithm cannot decide between several courses of action (e.g., Leeper et al. 2012).

Another strategy is to simply present the information operators need in a useful way that does not distract them from other tasks, such as allowing an operator to choose when a robot should be controlled (e.g., Glas et al. 2012), or enabling operators to make the final decision if a robot should investigate some area of interest (Daniel J Rea, Seo, et al. 2017). Both modelling and non-modelling approaches are used to enable some level of automation, freeing the operator from some tasks, but are designed to allow the operator to give input into when or how that automation takes place.

I note that robots or their interfaces can adapt physically to their operator. For example, a robot can also be aware of different users' physical abilities, and adapt its interface and controls to increase the accessibility of teleoperation to a variety of potential users (Balaguer et al. 2007).

In summary, situation awareness is a broad and important problem in teleoperation. Operators must build and maintain awareness of a robot's state and the environment around the robot to make effective decisions, and the robot should stay aware of the user to better understand how to interpret commands. Situation awareness is also a problem in video games, as argued in

Chapter 1, and have designed different interfaces to increase both player awareness of the virtual game world, and for the game to better understand the player's state and goals. Thus, this thesis investigates how video game interaction designs that make situation awareness easier for the player may be applied to teleoperation.

2.2.3 Interfaces for Controlling Robots

One of the main tasks that teleoperation researchers have worked to make easier has been how control robots. This may be adjusting the robot's body pose, such as moving a multi-jointed robot arm, or to help drive a robot through an environment. Control itself consists of many problems, including situation awareness, choosing levels of autonomy for an action, or dealing with physical problems like latency, described above. In general, controls need to be presented in a clear manner so that operators can understand and reason about how to command a robot to complete a task they may have – known as a gulf of execution that the operator must cross with the help of good interaction design (Norman 1986).

Semi-autonomous Robot Controls

Once an operator decides to move a robot, there are two general approaches: real-time or set-and-go controls. The latter is a semi-autonomous mode where some level of goal (destination, pose, action, etc.) is defined by the operator, and then the robot autonomously proceeds partially or completely to that goal (e.g., Singh et al. 2013; Quigley, Goodrich, and Beard 2004; Tsui et al. 2013). Once commands have been given, there is a delay while the robot proceeds, which an operator can use to deal with other tasks or robots (e.g., Glas et al. 2012; Olsen and Wood 2004).

This semi-autonomous operation style is beneficial as it enables operator multitasking, and long-term planning of robot actions (e.g., Liu et al. 2011; D. Sakamoto et al. 2009) see Figure 2.5. However algorithms may be imperfect and the real world can be dynamic, and so it may be necessary for the operator to provide more input during a task, such as help a kinematics simulator predict what position would be best to grip an option with a robotic arm (Leeper et al. 2012). Another potential drawback is that while attending other tasks, operators must maintain situation awareness of the teleoperation task, or reacquire it upon returning to the robot, potentially delaying task completion and adding workload to the operator (Donald,

Donald, and Thatcher 2015). Further, an operator may wish to edit or cancel an existing command in real-time, adding more complexity to the interaction.

The potential increased task performance with multitasking is desirable, but due to these tradeoffs, semi-autonomous teleoperation remains an active field of research. However, it is particularly suited to be used in multiple robot teleoperation scenarios where one operator controls multiple robots (e.g. Kolling, Nunnally, and Lewis 2012; Glas et al. 2012), as controlling multiple robots in real-time without automation is extremely difficult.

Control Strategies for Teleoperation

Even controlling a single robot is a challenging task that taxes an operator's cognitive resources (Steinfeld et al. 2006); seemingly basic tasks such as navigating a single, wheeled robot around a space are difficult enough that researchers have invented interfaces that aim to reduce the overhead required for a teleoperator in such a situation (e.g., Young et al. 2011; Barber et al. 2010; Singh et al. 2013). Specific strategies include following well-known control metaphors (e.g., a dog leash for a robot companion Young et al. 2011), visualizing the results of a command (Ashish Singh et al. 2013), using intuitive controls such as sketching paths in an image of the environment (Liu et al. 2011; Sugiura et al. 2010; D. Sakamoto et al. 2009), see Figure 2.5. What these works all share is a design rooted in familiar ways of acting and thinking

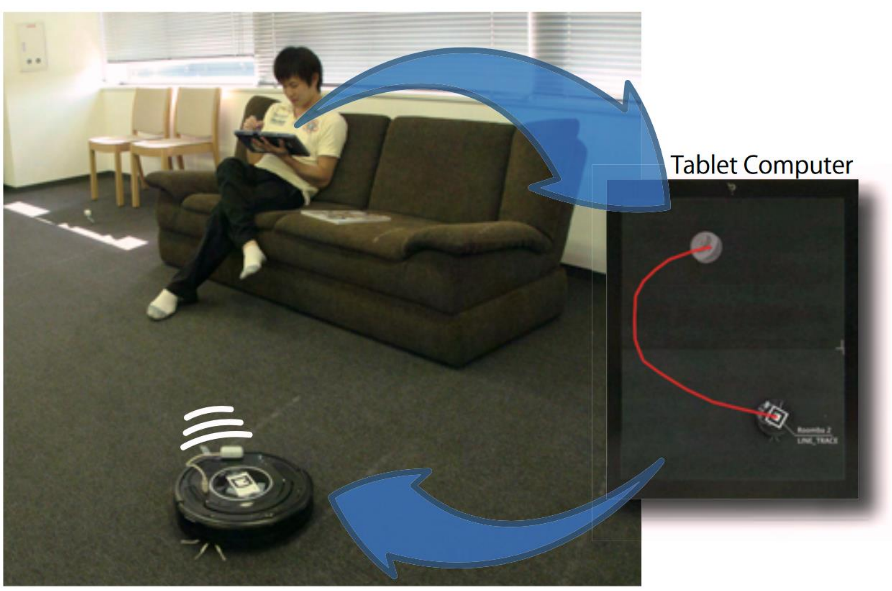

Figure 2.5. A sketch-based control interface for a robot that overlays commands in an overhead view of the real world to aid control and understandability of the robot's future actions. From D. Sakamoto et al. 2009.

(behavioral and cognitive psychology, human factors) instead of designing an interface around functional needs alone. The benefits of this strategy are improved operator performance and reduced operator workload. My work explores how video games use interfaces to solve similar control problems, and what benefits may be gained by applying their concepts to teleoperation.

Another factor that can affect operator workload is how fine-grained controls are. For example, an operator may need to define a precise path through an environment or grip an object at a certain angle; in these cases it is common to have complete control over robot movements with specialized interfaces designed for one robot's capabilities (e.g. D. Sakamoto et al. 2009; Glas et al. 2012; Hashimoto et al. 2011). However, complex controls can make some actions, especially common actions, tiring to manually perform repeatedly. For these situations, one strategy is to combine those common but complex commands into single actions that are easy to invoke (Barber et al. 2010). By understanding the tasks operators wish to complete with a robot, the interfaces can be made more manual or more automated to ease teleoperation.

Some design strategies are more goal-oriented and attempt to improve specific performance metrics or create certain user experiences. Many broad approaches have been adopted, including control methods to reduce task completion time or collisions (e.g., Daniel J Rea, Hanzaki, et al. 2017; Leeper et al. 2012), awareness of the remote area (e.g. Nielsen, Goodrich, and Ricks 2007; Singh et al. 2013; Seo et al. 2017; Endsley 2016), and mental resource management to improve overall operator performance (e.g. Hacinecipoglu, Konukseven, and Koku 2013; Chen, Barnes, and Harper-Sciarini 2011). My approach directly builds on these results by exploring new interaction designs to improve teleoperation that borrow inspiration from techniques in video games.

The common trend in all these projects is that control systems should be designed to help operators understand what commands are available and how the command will affect the robot and the environment (such as by having the commands be input on an image of the real world space), and that control interfaces should take into account the limitations of an operator (e.g. the difficulty of selecting and commanding multiple robots). Both of these lessons relate to the concept of embodiment, described earlier.

Telepresence

Telepresence is the study of an operator's sense of being physically present in the remote environment (Kristoffersson, Coradeschi, and Loutfi 2013), often to facilitate human-human interaction through a robotic proxy (e.g., Tanaka et al. 2016; Wentzel et al. 2015; Radmard, Moon, and Croft 2015; Tsui et al. 2013). A sense of presence has been linked to a variety of benefits including increased understanding of spatial tasks and relations (Kulik 2009; Stoakley, Conway, and Pausch 1995), more engagement (Schuemie et al. 2001), and better memory (Schoor 2006; Max M., Sarah M., and Joseph R. 1998). Presence has also been linked to more enjoyment (Schoor 2006; Tsai, Shen, and Fan 2014), and thus is also seen as an important design goal in video games (Jennett et al. 2008; Mekler et al. 2014). In this thesis, I do not specifically investigate presence or telepresence scenarios, however as presence is often a goal in game design, my approach and results may indicate that video game interfaces could improve telepresence applications as well.

Similarities between Driving Vehicles and Teleoperation

Especially for navigation tasks, teleoperation is similar to driving a vehicle: a person wants to take a vehicle from some origin to a destination quickly without causing accidents. Supporting the results in interaction design (Section 1), research in traffic psychology has shown that a driver's psychological state can change their driving behavior (J.A. Groeger and Rothengatter 1998; John A. Groeger 2002). These changes may be due to the perception of the vehicle itself (e.g., sporty vehicles, Eyssartier, Meineri, and Gueguen 2017), the surrounding environment (Michon 1985), the driver's mood (Precht, Keinath, and Krems 2017), or even the physical controls of the vehicle (Blommer et al. 2017; McIlroy, Stanton, and Godwin 2017). I extend some of this work to teleoperation by investigating if I can use affective feedback to change an operator's mental state, and therefore change their driving behavior (Chapter Five). I further investigate changing a driver's perception of the robot they drive, and see how that impacts an operator's mental state and driving ability (Chapter Four).

2.2.4 Summary: Teleoperation Is Difficult

In general, teleoperation work focuses on conveying the state of the remote world to the operator and enabling the operator to control the robot quickly and easily. However, while extensive literature in human-robot interaction continues to grow, teleoperation remains difficult due to challenging and multifaceted problems.

My work builds on previous human-robot interaction research by surveying the interface techniques in video games and building a taxonomy of some video game solutions to awareness and control problems (Chapter Six). By comparing the previous work in the teleoperation to my taxonomy, we see trends pointing to the applicability of video game interface techniques to robot teleoperation. I further add to the field by suggesting extensions to current research directions based on promising video game interfaces I observed. Based on some inspiration from video game interfaces, I further attempt to combine elements from social human-robot interaction, and teleoperation, formalizing new approaches in teleoperation interaction design.

2.2.5 Social Human-Robot Interaction

Robots that are physically close to people, unlike teleoperated robots that are separated, which makes it natural the robot to use human-like and social language and behaviours to better interact with people. While autonomous robots are just computers with a body that can move through physical space, one may expect little relevance of social abilities and effects when people interact with them. However, there is now a large body of evidence that, similar to the computers-as-social-actors paradigm (CASA, Lee and Nass 2010), people react to and interact with robots as if they were social beings, and that robots are fundamentally treated in a more social way than other technology (Breazeal 2003; J. E. Young et al. 2009; J. E. Young 2010). This is often seen as an anthropomorphic or zoomorphic effect – treating robots as if they are humans or animals respectively (Bartneck et al. 2009; J. E. Young 2010; Ashish Singh and Young 2013). Why this happens is still open to exploration, but it is thought to occur in part due to the embodiments of people and robots (Section 2.1.5, Klemmer, Hartmann, and Takayama 2006). This is further evidence of how interaction with robots should be treated separately from other technology (J. E. Young 2010; Breazeal 2003).

Socially embodied interaction works for both robot control and for providing feedback to people around the robot. For example, a person could command a robot to help them with a task by using natural language and body language such as pointing and gesturing with their hands and head (Breazeal et al. 2005). Similarly, people can understand feedback from a nearby robot presented in embodiments we are familiar with, such as robots using gaze, facial expressions, or natural language (Edsinger and Kemp 2007; Gleeson et al. 2013; Admoni and Scassellati 2017). Not only can a person interact with a robot using social communication, people will also think about robots socially and can form social connections with a robot that can give people strong emotional responses (Sung et al. 2007; Seo et al. 2015b). At this point, the model of operating a robot (giving commands and receiving feedback) needs to be expanded to consider the social component of interaction and the new interaction opportunities it creates outside of traditional operation. This field is known as social human-robot interaction.

Robots can use social skills to improve or influence their interactions with people. For example, they can use body language to communicate (Breazeal et al. 2005; Gleeson et al. 2013), use visible hesitation and verbal stalling techniques while "thinking" (waiting for processing to complete, Moon et al. 2013; Ohshima et al. 2015; Glas et al. 2012), adapt their behaviors to be more compelling or attractive to their human interaction partner (Banh et al. 2015; Sanoubari et al. 2019), or even purposefully control and influence group dynamics between humans (Jung, Martelaro, and Hinds 2015; D. Sakamoto and Ono 2006). Robots can even affect people emotionally by their actions (Riek, Paul, and Robinson 2010) or presence (Shibata and Wada 2011; Sung et al. 2007).

People intuitively, and, perhaps subconsciously, respond to robots as if they were intelligent social agents regardless of whether the robot has been programmed to be social or not (Sung et al. 2007; Forlizzi 2007). For example, people can build empathy for robots who work and talk with them (Seo et al. 2015b; P. H. Kahn et al. 2012; Riek, Paul, and Robinson 2010), have their relationships with other people changed by a robot's actions (Jung, Martelaro, and Hinds 2015; D. Sakamoto and Ono 2006), or get angry at robots they believe are breaking rules or social conventions (Short et al. 2010; P. H. J. Kahn et al. 2012). People will even keep a robot's secrets from their owners (P. H. Kahn et al. 2015; Seo et al. 2015a) or name and buy clothes for a robot that has no programmed or explicitly designed social elements (Sung et al. 2007). Even further, robots can pressure people to do more work than they are willing through use of a position of authority in the social hierarchy (Geiskkovitch et al. 2016; Cormier et al. 2013).

People report understanding that robots are just computers and are not conscious beings (Sung et al. 2007), but they still respond as if robots were just that.

One of the important lessons learned from social human-robot interaction is that people respond more socially to robots than they do with other technology. The field itself also echoes the computers-as-social-actors idea that people reflexively respond to technology in social ways, and that we can further leverage this by designing technology, or in this case robots, with human social tendencies in mind. As touched upon above, social and affective psychology are considerations in interaction design related to how people think.

An open question remains as to how robot teleoperation interface designers can harness and leverage these social and emotional interaction paradigms. However, video games use social techniques to affect the player, such as using companions as part of the user interface [The Legend of Zelda: Ocarina of Time, Nintendo, 1998; Doom, id Software, 1993], or building rapport between the player and other computer characters to later leverage in heated debates [such as in Mass Effect 2: BioWare, 2010]. In this thesis, I explicitly pursue and establish this link, and investigate how social techniques can affect teleoperator behaviour and experience (Chapter Five).

2.3 Evaluating Human-Robot Interaction

The problems of teleoperation, touched upon above, can be seen as interaction problems. This enables us to use human-computer interaction and interaction design evaluation techniques, such as user studies, statistical analysis, and exploratory qualitative methods. The purpose of these methods, which aligns with my goal, is to learn about people to inform the design of future technology (J. E. Young et al. 2010; Preece, Sharp, and Rogers 2015)

I can evaluate interaction designs, including those in teleoperation, from multiple angles, such as task performance, human factors (comfort, workload, etc.), cognitive requirements, and social and emotional effects. Experiments can be performed in the field or in a lab with a mock-task to simulate real use. Field studies have the advantage of being very ecologically valid – observations capture real people using technology to solve real problems. However, real life can vary and be unpredictable, thus lab studies can provide control and consistency, reducing noise. As robots are still emerging in the world, and as many current uses of teleoperation are for emergency situations, it is common to use mock-situations relevant to the intended application, such as a mock search-and-recue scenario (e.g., Seo et al. 2017; Drury, Scholtz, and Yanco 2003), a navigation task for driving interfaces (e.g., Hacinecipoglu, Konukseven, and Koku 2013; Daniel J Rea, Hanzaki, et al. 2017; Daniel J Rea and Young 2018), or a social interaction via a robot (e.g., telepresence, Rae, Takayama, and Mutlu 2013; Tsui et al. 2013). I briefly discuss methods pertinent for this thesis and my target applications, and refer other readers to an excellent and still relevant review by Steinfeld et al (Steinfeld et al. 2006).

Workload is an important measure and consideration for the challenge felt during teleoperation (Chen, Barnes, and Harper-Sciarini 2011; Chen, Haas, and Barnes 2007). As discussed, workload itself is not necessarily correlated with task performance, but reduces cognitive resources over time, and any task that requires mental abilities that exceeds a person's current ability will likely result in negative performance impacts. Relatedly, teleoperation researchers have noted the importance of reducing the number of interactions, attention demand, and memory requirements for interaction (M.A. Goodrich and Olsen 2003). Due to this link, the teleoperation field does consider lowering workload a valuable impact of an interface (Chen, Haas, and Barnes 2007; Steinfeld et al. 2006). In my work, I measure, and in one case specifically target, impacts on operator workload.

Robots often have functional goals and thus it is important to measure if those goals have been achieved, how efficiently they were achieved, and any other impact to the robot's surroundings. Typically, tasks will have a completion time measured, with the assumption that faster completion time suggests a better interaction design (e.g., Radmard, Moon, and Croft 2015; Chen, Haas, and Barnes 2007; Daniel J Rea, Hanzaki, et al. 2017). Others may instead (or additionally) record number of completed tasks (if repeated, such as in Daniel J Rea, Seo, et al. 2017), or develop a more complex model of what it means to succeed at a teleoperation task (e.g., Seo et al. 2017). As teleoperation is difficult, mistakes are common, even for experienced operators (Drury, Scholtz, and Yanco 2003; Yanco and Drury 2004a). And so, measures of incidents, such as collisions with the environment or people, are often counted (e.g., Cotter 2014; Weibel and Hansman 2005; Chen, Haas, and Barnes 2007; Daniel J Rea and Young

2018). Thus, in addition to workload, collisions and completion time are both common measures throughout my interaction evaluations.

When relevant, it may be useful to observe potential emotional or social effects on the operator or those who interact with the operator's robot. As with cognitive workload, there are potential physiological metrics that can be measured to deduce a person's emotional state, but these can be difficult to measure and may be confounded with other physical responses. Thus, self-report questionnaires, such as the Self-Assessment Manikin, can be administered to gain insight into emotional state (Morris 1995), and have been used successfully in social human-robot interaction (such as in Singh and Young 2013; Sharma et al. 2013; Thiessen et al. 2019). Free-form feedback in writing or interviews can also provide insight into emotions and social impact, as well as a more insight into the thoughts and reasons behind a person's actions, but often require qualitative analysis (A. L. Strauss 1987).

Interaction design is a multifaceted discipline (Section 1), and so I take appropriately multifaceted evaluation measures. I include observations and analysis of objective measures such as completion time and number of collisions to understand to understand the performance impacts of my designs on operators. I combine self-report questionnaires with qualitative analysis of open-ended feedback to better understand teleoperation experience, to gain insight into *why* performance impacts occurred, and what psychological influences my interaction design may have had.

2.4 Video Game-inspired Interaction

The general link between interaction design and video games is well established. Game developers have leveraged human-computer interaction techniques such as iterative design (Burgess 2014), user experience and user-centered design (D. Sakamoto 2015), and even lower-level interface techniques such as layout and cursor management (Candland 2016). Game developers have also experimented with and helped popularize novel interface technologies such as motion tracking (e.g., WiiMote or Kinect video game hardware peripherals), and touch screens (such as the Nintendo DS handheld console). Because of the success of these game interfaces, I argue that the investigation of using game techniques to improve teleoperation interfaces is an important approach.

2.4.1 Game Design and Gamification

Video game-inspired designs have already been applied in the broader field of human-computer interaction. For example, *gamification* uses game design to improve emotional engagement (Langer, Hancock, and Scott 2015) or ability to learn software (Malacria et al. 2013; Li, Grossman, and Fitzmaurice 2012) and robot systems (Labonte, Boissy, and Michaud 2010). Game-like virtual worlds have been used to explore ideas that would be difficult to implement physically (Atkinson and Clark 2014; Feldmaier, Stimpfl, and Diepold 2017), and games can be used to generate data to train robot artificial intelligences (Walther-franks et al. 2015). While these works demonstrate the usefulness of video-game inspired interfaces in software in general, I further posit that they can be successful on a larger scale – and may be even more relevant – in teleoperation, where the problems of remote control and controlling a virtual character share many similarities.

Gamification is the use of game-based design, game elements, and game characteristics in non-game contexts (Deterding et al. 2011). The general focus of gamification research and applications, however, has focused on boosting incentive, engagement, and user motivation in software (e.g. Deterding 2012; Guo et al. 2012; Malacria et al. 2013; Hamari, Koivisto, and Sarsa 2014; Kirman et al. 2013; Li, Grossman, and Fitzmaurice 2012). I complement this research direction by exploring how other game design concepts can be used as inspiration for new robotic teleoperation interfaces. It has been noted (in Deterding et al. 2011) that game design elements for gamification could include: game interface design patterns, game mechanics design patterns, game design principles and heuristics, game models, and game design methods, as well as input devices.

Research into games themselves has investigated how interaction design in the game can affect player performance and experience (Mekler et al. 2014; Ermi and Mäyrä 2005; Hochleitner et al. 2015). For example, information presentation inside in-game heads-up-displays has been connected to how immersed a player becomes in game (Iacovides et al. 2015; Babu 2012; Caroux and Isbister 2016). How games design for challenge and mastery of the game's interface can increase engagement and enjoyment (Nah et al. 2014). Interface layout choices such as integrating interfaces into the game world (Ogier and Buchan 2017), or designing to affect user emotions (Johnson and Wiles 2003) also have improved performance and experience in games. These projects showcase the links between video games and interaction design and experience, and this thesis explores how the interactions designed in games could promote similar performance or experiential benefits in teleoperation.

2.4.2 Video Games and Teleoperation

Teleoperation reasearchers have developed and evaluated successful interfaces that share similarities with video game interfaces. For example, researchers have used techniques also similar to those used in game interfaces to improve teleoperation camera use (Saakes et al. 2013; Hashimoto et al. 2011; Keyes et al. 2006; Seo et al. 2017), or improve robot controls (D. Sakamoto et al. 2009; Kolling, Nunnally, and Lewis 2012). Such research demonstrates that interfaces with elements common in video games can improve teleoperation and serve as initial evidence and encouragement for my approach of formally investigating video game interfaces. I add to this body of work with my own evaluated designs, using video game inspiration to design teleoperation interfaces that help operators maintain awareness (Chapter Three - Daniel J Rea, Seo, et al. 2017), or change operator perceptions of a robot (Chapter Four - Daniel J Rea and Young 2018).

Richer *et al.* proposed the Video Game Based Framework in order to analyze interfaces for human-robot interaction (Richer and Drury 2006). They built a framework to characterize and classify robot control interfaces and their components with video game vocabulary focusing on game control: controllers, camera control, how input is used to generate commands, etc., and they note that the interface design aspect of video games is broad and out of scope in their work. I extend this initial exploration into video game design by providing a broader, data-driven, and formal analysis of video game interfaces across a defined sample of games that creates a structured understanding of video game interface techniques, the messages they convey, and, analyses of how they convey them. In particular, I focus on interface techniques that specifically leverage psychological effects, which was an area not covered in depth in Richer *et al.*'s work. My approach is also tangential in purpose: I wish to create a taxonomy of game techniques to explore how video game interface concepts could be applied human-robot interaction and help identify techniques to inspire new research.

Input hardware common in video games has also been shown to improve teleoperation. Input devices can provide a user with ways to think and reason about what they can do with a system (Maier and Fadel 2009), and video game hardware has often experimented with input devices to enhance players' gaming experiences. These include the modern, many-button and multi-joystick controllers, and unique controllers (rhythm controllers, haptic feedback controllers, vision-based controllers such as the Kinect, the Wii Fit Board). Research has noted that modern controller designs can increase the accessibility of teleoperating the robot (Singer 2009), including more standard designs (sometimes even included with commercial robots ("IRobot 510 PackBot" 2015; "Clearpath Robotics: Controllers" 2017), and less common controller designs (e.g. Wii Remote – Young et al. 2011, or the Microsoft Kinect – Levy-Tzedek et al. 2017). These successes should encourage the investigation of other types of game hardware, as well as the interfaces that work with them, in the context of teleoperation.

I propose to move beyond a few targeted video game-inspired interface examples: learning from video games could be an overarching approach for improving teleoperation, and video games should be a common source of design material and inspiration. Thus, I explicitly and systematically examine current game interaction techniques, as well as the successes in the above research.

2.5 Related Work - Summary

Teleoperation itself is at an intersection of interaction design, human-computer interaction and robotics; the interaction design and software usability literature is important and applicable to teleoperation, however the remote operation of a robot adds new challenges that must be considered, such as maintaining situation awareness and managing physical issues such as unreliable sensors. Interfaces aim to help improve the teleoperation experience by reducing an operator's mental load and improve operator performance (i.e., completing tasks safely and efficiently). Modern teleoperation solutions create new interfaces to improve controls, information displays, and have drawn from psychology, sociology and traffic safety research. My work aims to build upon previous improvements in teleoperation interfaces by exploring how video game interfaces may leverage human factors and psychological tendencies in humans to provide further benefits to the operator's performance or experience.

I add to the teleoperation literature in a number of ways. My approach is based in interaction design, and showcases the benefits of considering the mental state and processes of operators in my designs. This extends to a demonstration across all projects on the importance of user experience in teleoperation design, and how we can design to improve the teleoperation experience which can, in turn, improve operator task performance (Chapters Three – Five). Specifically, I demonstrate how considering perception and cognitive processes can lead to improved visual interface design (Chapter Three), and how we can design to influence an operator's perception of a robot's abilities and performance by targeting the operator's mental models of the robot (Chapter Four). I further provide the first bridge and evidence for the application of social human-robot interaction techniques to teleoperation interface design (Chapter Five). My framework takes a broader view of video game interaction design, and provides evidence for how video games themselves leverage these interaction design principles to improve user experience. I further use my framework to better understand and highlight the similarities of the problem and design spaces of video game and teleoperation interaction design (Chapter Six), demonstrating that video game-inspired interfaces should be an approach further explored by the community.

Video game interfaces have been demonstrated to add benefits when added to traditional software, primarily focusing on engagement with software, and learning software. I argued that video games and teleoperation share many similarities, and, indeed, preliminary work in teleoperation has found benefits to the initial application of interfaces similar to video game interfaces. While my search found video-game inspired research is an emerging trend, there was a lack of overarching strategy in the literature for using game elements in telerobotics interfaces, which I aim to contribute to with my research. Thus, my work aims to extend teleoperation, gamification, and interaction design in general with a broader and systematic evaluation of applying lessons from video game interfaces to teleoperation.

Chapter Three: Directing Attention during Teleoperation

One problem faced by both video game developers and teleoperation interface designers is to draw attention to various areas and events in the remote or virtual world. Specifically, if a player or teleoperator is moving through an environment, the system may sense something of interest that the user should investigate in more detail. In other words, it should direct the attention of the user to that area of interest. In this chapter, I present the results from an iterative design process that I conducted to develop different attention drawing interfaces for teleoperation. My interfaces are inspired by video games, grounded in the psychology of perception and attention, and evaluated in a teleoperation scenario.

I evaluate my designs with a three-cycle iterative design process, evaluating interfaces on their attention-grabbing capability and impact on operator performance. My results show that operators perform poorly without attention-drawing interfaces, and that all of my interface designs improve operator performance and reduce cognitive load compared to no aid at all. Further, I synthesize the results from my multi-stage process and detail which design parameters impact operator cognitive load and task performance. Specifically, full-screen interfaces can lower cognitive load, but can increase response time when attention is drawn, and animated cues may improve number of regions of interest found but increase cognitive load. The results of my work are formally tested, and theoretically grounded cues for attention drawing in teleoperation, and my positive performance improvements from video game-inspired attention drawing interfaces provides my initial support for this thesis. Further, our approach demonstrates that an iterative design approach may be useful in adapting game techniques to apply them to teleoperation for further benefits to the field.

Parts of this chapter have been taken in part or in full from the following publication:

Daniel J. Rea, Stela H. Seo, Neil Bruce, James E. Young. "Movers, Shakers, and Those Who Stand Still: Visual attention-grabbing techniques in robot teleoperation." In Proceedings of the ACM/IEE International Conference on Human-Robot Interaction. ACM/IEEE. 2017. (24% acceptance rate)

3.1 Attention, Teleoperation, and Video Games

Teleoperation interface improvements have aimed to increase operator performance for applications including the military (Richer and Drury 2006), industrial (Ashish Singh et al. 2013) or domestic tasks (Labonte, Boissy, and Michaud 2010; Mast et al. 2015), and remains an ongoing research challenge. One such goal of this improved efficiency is to enable fewer people to control or monitor more robots, increasing the human-robot ratio (Yanco and Drury 2004b; Olsen and Wood 2004), and getting more work done faster. However, increasing the information given to operators, such as simultaneous video feeds from separate cameras or robots, results in higher cognitive load and operator error (see Chapter Two, Section 2.1.2, and Drury, Scholtz, and Yanco 2003; Singh et al. 2013; Radmard, Moon, and Croft 2015; Pylyshyn and Storm 1988). As such, a primary goal of teleoperation usability research is to improve overall operator effectiveness: provide operators with the tools and information they need to perform their tasks, without overloading them mentally. I follow this theme in this project, exploring tools to increase performance in teleoperation.

One common task in teleoperation is visual search: in urban search and rescue, operators may be searching for damaged buildings or injured people, unmanned aerial vehicles may be searching for intruders in an area, or a person operating a telepresence robot may be searching for a face in a crowd at a conference. Visual search tasks can be made easier by providing assistance with emerging computer vision techniques (e.g. Teng, Kuo, and Tara 2013; Bruce, Catton, and Janjic 2016) by automatically identifying potential points of interest, and indicating (*cueing*) the points for inspection. However, such attention cues – the interface technique used, such as graphics or sound, to indicate to the operator – cannot be too subtle: while their attention is focused people may not notice events outside their immediate focused area (Treue and Martínez Trujillo 1999). On the other hand, attention cues that are too intrusive can be annoying, frustrating, and distracting to the point of negatively impacting the continued search (Tasse, Ankolekar, and Hailpern 2016). Thus, while indicating areas of interest could be helpful in teleoperation, it is not clear how this information should be cued to an operator to effectively gain their attention, without overly distracting from other tasks.

Similar to teleoperation search tasks, video game players are often also visually searching for items to collect, incoming enemies, their next navigational objective, and more. Games have experimented with numerous methods to point out important things, such as AI companions that highlight enemies or important parts of the environment [e.g., The Legend of Zelda: Ocarina of Time, Nintendo, 1998; Mass Effect 2, BioWare, 2010], automatically making interesting areas more visually salient [e.g., Super Mario Galaxy, Nintendo, 2007; Grand Theft auto IV, Rockstar Games, 2008], or employ methods to bring the player's attention to an incoming event [e.g., Goldeneye 007, Rare, 1997; Bioshock, 2K Games, 2007; Doom, id Software, 1993]. I take inspiration from two common video game interface elements: the targeting indicator (Figure 3.1, right), and the damage indicator (Figure 3.1, left).

Damage indicators in video games serve a simple purpose: to let players know their character has been injured. This may be a general damage indicator, simply indicating damage occurred (Figure 3.1, left), or an indicator that also indicates where the damage came from (where the source is [e.g., Overwatch Blizzard, 2016]. Targeting indicators are used to display what a player's actions are (or should) be focusing on (Figure 3.1, right). As players may often switch controls or strategies depending on if they are targeting something, indicating what and where they are targeting is very important. In this sense, both targeting and damage indicators are used to indicate an important event, location, or item that the player should pay attention to. I take design inspiration from these indicators to create new interfaces for teleoperation that direct attention to important areas that appear in a robot's video feed. Further, to ground my exploration into adapting these gaming techniques, I leverage related work on visual attention in psychology to help ground and evaluate different visual attention guides.

Figure 3.1. Two examples of video game attention drawing indicators. Left: The entire screen flashes white and large health bars appear on the screen as the player is shot. This important event uses large full-screen, salient designs [Goldeneye 007, Rare, 1997]. Right: A player targets an enemy, and bright yellow triangles appear at the edge of the screen, move towards the plant enemy, and begin circling it. The player's eyes are drawn to this movement, and the bright cue around the targeted enemy makes it easy to spot [The Legend of Zelda; Ocarina of Time, Nintendo 1998].

We conducted an iterative design exploration of using visual cues in a multi-robot search and rescue context. I iteratively designed, implemented, and evaluated cue variants starting with inspiration from video games, such as above, while basing my design choices on the literature on the psychology of attention and perception. This background work (detailed in Section 3.3, below) describes how the size of an interface element, as well as the element's movement, can grab attention. Thus, I designed my cues based on cue proximity to the target (e.g., a small interface located right at the target to encode its location versus a full-screen interface to gain attention), and cue motion (moving to gain attention or static to not be too distracting).

I performed one pilot experiment and two design iterations that included evaluating my designs and their effects on operator visual-search accuracy, response time, and cognitive load. Results indicated that moving cues can help operators find more lights than static cues. Further, cues located at a light, particularly when moving, can be the quickest for operators to assess, but also increase cognitive demand. Well-designed full-screen cues can achieve similar task effectiveness while simultaneously lowering the cognitive load required on operators. In addition to my study results and reflection on these parameters, I present a set of tested, iteratively designed cues that aide operators in visual search tasks without negatively impacting operator cognitive load.

Figure 3.2. Four camera feeds from four different robots are shown, where an operator needs to monitor for points of interest (bright green lights). Visually cueing a point of interest for an operator observing four robots searching an area. This example uses a bouncing circle that draws attention to the point of interest (green light, red lines indicate circle movement, and are not shown in the interface).

This work illustrates how video game interface design techniques can be leveraged to adapted to teleoperation to help direct an operator's attention during a visual search task. While I focused on performance in attention and reductions in cognitive load, video games use other interfaces to help in this regard as well, such as for supporting player state awareness, navigation, or understanding the virtual environment. This chapter's overall approach – grounding techniques in perception literature and iteratively designing and evaluating interfaces – may be applied again to these new interfaces to bring them to teleoperation, perhaps further reducing cognitive load and improving an operator's use of attention resources in a variety of teleoperation situations.

3.2 Interfaces to Reduce Cognitive Load and Direct Attention

Supporting teleoperation (aiming to increase performance and lower cognitive demands) by modifying how video feeds are presented is an active research area (Drury, Scholtz, and Yanco

2003; Hashimoto et al. 2011; Richer and Drury 2006; Labonte, Boissy, and Michaud 2010). Much of this has revolved around camera location and viewpoint choice, for example, egocentric views enable an operator to see the world from the robot's perspective (Labonte, Boissy, and Michaud 2010; Kristoffersson, Coradeschi, and Loutfi 2013; Drury, Scholtz, and Yanco 2003), environmental views provide robot-in-context information (Hashimoto et al. 2011), while bird's-eye and third-person views provide this context calibrated around the robot as it moves (Fukatsu et al. 1998; D. Sakamoto et al. 2009; Saakes et al. 2013). When multiple feeds are displayed, research has explored layout options, such as displaying all screens in a tiled fashion to maximize available information (Stoica, Salvioli, and Flowers 2014), or using picture-in-picture techniques to prioritize screen real-estate toward more important views (Krajník et al. 2011). Feeds can be hidden until requested (Glas et al. 2012), with operators perhaps rotating through them (Atrey, Hossain, and El Saddik 2008) – reducing information load and saving screen space (Glas et al. 2012). Rather than projecting views or representing camera placement, I aim to support teleoperation by providing task-specific help: drawing operator attention to points of interest while mitigating the negative effects of the interruption.

The video feed itself is commonly augmented with graphics to represent relevant information, similar to video game designs (e.g. Figure 3.1). For example, the operator can be presented with different types of sensor data (Richer and Drury 2006; Keyes et al. 2006), notifications (Chen, Barnes, and Harper-Sciarini 2011), or task information with ecological design (Ashish Singh et al. 2013; Hashimoto et al. 2011; Mast et al. 2015). I expand on this work by investigating how to notify users of potential points of interest within a video feed.

In computer vision, saliency detection refers to the problem of detecting image regions that are likely to be salient to a human viewer. This is a complex problem which considers physiology of vision as well as psychology. Saliency detection has been used successfully to shift (Veas et al. 2010) and predict (Bruce, Catton, and Janjic 2016; Torralba et al. 2006) gaze, improve visibility in augmented reality (Kalkofen et al. 2013), and model how people process visual context (Chun 2000). It has also been used to modify video scenes, and to draw human attention toward objects (e.g. Veas et al. 2010); in this case, the changes were subtle and the goal was for impacting viewer tendency to look at areas, not direct attention to immediate concerns. For teleoperation, saliency detection has been applied only rarely, to inform interface design of sensor readouts (Chen, Barnes, and Harper-Sciarini 2011), or to minimize network usage for visual data (Teng, Kuo, and Tara 2013); here, low-saliency regions use lower resolution, effectively blurring them. I extend this work by investigating how visual cues can be designed to appropriately draw operator attention to such points identified through saliency detection, while aiming to lower operator distraction that may impact task performance.

Human visual attention, particularly for visual search, has a rich history in psychology. Studies frequently focus on static abstract images (Chun 2000; Klein et al. 1999; Kristjánsson, Jóhannesson, and Thornton 2014; Thornton et al. 2014), or abstract videos with colors or shapes changing and moving against solid backgrounds (Abrams and Christ 2003; Atrey, Hossain, and El Saddik 2008; Burke et al. 2005; Kristjánsson, Jóhannesson, and Thornton 2014; Klein et al. 1999; Mack and Rock 1999). Studies with natural scenes tend to use static images (Treue and Martínez Trujillo 1999; Torralba et al. 2006; Chun 2000). Some video work in natural scenes focuses on closed-circuit television monitors (Atrey, Hossain, and El Saddik 2008; Donald, Donald, and Thatcher 2015; Howard et al. 2009; Stainer, Scott-Brown, and Tatler 2013). Robots, unlike CCTV cameras, move throughout their environment freely, presenting highly dynamic and noisy camera views from constantly shifting perspectives; further, increasingly dynamic environments tax users' attention resources (Pascoe, Ryan, and Morse 2000; Tarasewich 2003). As such, prior attention results must be specifically evaluated in the unique, high-demand teleoperation context.

Notification work for general desktop applications has focused on *when* to draw attention (McCrickard and Chewar 2003; Horvitz, Jacobs, and Hovel 1999; Horvitz et al. 2003; Bailey and Iqbal 2008), which does not apply to my task, where operators must be immediately notified and respond in a short time (before a target leaves the screen). Work on *how* to draw attention is much more limited. Some has highlighted how interruptions can be distracting or annoying, and has investigated how to minimize these problems (Veas et al. 2010), while (recently) noting the lack of solutions to this problem (Tasse, Ankolekar, and Hailpern 2016). Further, heavy use of interruptions can lead to users ignoring them (Burke et al. 2005). These results, primarily from web and desktop applications, motivate the need for my research in the visually-intense and noisy teleoperation task, exploring attention-drawing cue design that balances being attention-grabbing while not being distracting or being ignored due to fatigue.

3.3 Attention and Perception

Human visual attention – how people choose what to focus their vision resources on – is a well-studied topic in biology, neurology, psychology, etc. Attention can be defined as an enhanced response to stimuli at an attended location and, as a result, reduced response to stimuli elsewhere (Treue and Martínez Trujillo 1999). Thus, I can expect people to have increased focus on some task elements (e.g., during searching, driving, reading), and, inversely, difficulty noticing things outside of their focus (Rensink, O'Regan, and Clark 1996), even highly-salient points of interest – this is called inattentional blindness (Mack and Rock 1999). This is especially difficult during noisy dynamic tasks, such as teleoperation (Simons and Chabris 1999). I aim to work within human patterns of attention to devise visual mechanisms to help gain people's attention and direct it to points of interest, with minimal overall hindrance or additional strain on cognitive resources.

One technique for focusing attention, called goal-based attention, cognitively directs attention to known criteria or stimuli, such as a known suspect on CCTV (Howard et al. 2009). Goal-based attention is relevant to teleoperation as operators often have specific, if broadly defined, tasks that drive visual search and cognition such as "find and rescue all victims in a disaster." Complicated search goals (multiple criteria, complex shapes) and environments, such as disaster environments, reduce the effectiveness of goal-based attention (Howard et al. 2009; Rensink, O'Regan, and Clark 1996; VanMarle and Scholl 2003). Increasing the number of cameras will also reduce the effectiveness of goal-based attention due to the increased search area (Stainer, Scott-Brown, and Tatler 2013). Goal-based attention quickly reaches limitations in complicated tasks and environments that may be present in tele-robotic search and rescue.

Alternatively, stimulus-driven attention draws a person's attention to salient stimuli, such as bright lights, motion, or high contrast graphics. Interfaces could have objects appearing (Franconeri and Simons 2005), elements starting to move (Abrams and Christ 2003; Abrams and Christ 2006), or motion perpendicular to other motion in the visual field (Franconeri and Simons 2005); not all changes are similarly salient, for example, color shift, or motion types such as receding motion or movement parallel to other ongoing motions, have been found to be less effective at drawing attention (Franconeri and Simons 2005; Abrams and Christ 2003).

Stimulus-driven design suffers less from fatigue in comparison to goal-based attention, and further suffers less from inattentional blindness, important for long-term attention (vigilance, Donald, Donald, and Thatcher 2015). As such, I design cues leveraging stimulus-based attention to mitigate some of the limitations incurred by the operator's goal-driven attention.

3.4 Cue Design Process

My investigation into how cue design impacts teleoperator performance employed an iterative design process: I drew from perception and attention literature to inform design, devised a mock urban search and rescue task for evaluation, implemented my cues into the mock task, and conducted formal experiments to learn of the impact of my cue designs. My results informed the design of new cues and conducted more experiments, for a total of one pilot (9 participants) and two formal studies (with 20 participants each).

3.4.1 Cue Evaluation Test Bed

I developed a test bed that engages participants in mock urban search and rescue, performing visual search on teleoperation feeds, enabling me to test the impact of my cues on visual search.

Task

Participants monitored a collage of four tiled video feeds from teleoperated robots exploring a mock-disaster environment (Figure 3.2) and were asked to search for stimuli that represented points of interest (e.g. potential victims, dangerous equipment). Participants tapped on a touch-sensitive screen near the stimuli to show they had identified it.

For my stimulus, I aimed for an abstract stimulus that would more readily generalize to a broad range of teleoperation tasks. As such, I avoided being domain-specific, and potentially confounding variables such as shape or pattern. My design goal was for an abstract, generalizable stimulus which is unambiguous once found, yet still difficult to find. I chose bright green point lights as my target. Further, I aimed to increase visual search validity by using a visually noisy scene with realistic robot movement and video quality, building on existing fully abstract perception work by investigating attention in a more representative visual environment.

Teleoperation Videos

I pre-recorded my robot teleoperation videos for consistency across participants – I could control when, where, and how often a participant saw a light. As participants only monitored the robot's camera feeds, and did not actually see the teleoperators, this is equivalent to live operation for my evaluation purposes.

I arranged a room to have furniture, electronics, and debris scattered around (Figure 3.2), and remotely controlled a NAO H25 robot over Wi-Fi traversing the space. The video was recorded from the robot's head camera (640x480 at 17 FPS).

I recorded five videos, each having a unique room and debris arrangement, while maintaining similar visual clutter, layout, lighting conditions, and robot movement properties (speed, frequent turns, minimal stopping). I compiled five different (but comparable in character) four-tile collages for a repeated-measures study design (Figure 3.3). I modified the video selection and position in the collage using incomplete Latin Squares to minimize learning effects. Each video and collage lasted six minutes and four seconds long.

Stimuli (light) Placement and Timing

I placed several centrally-controlled green LED lights throughout the mock environment to serve as my stimuli. Light timing and placement posed several challenges. First, only one light at a time should be illuminated in the entire collage, to avoid confusion over which stimulus a participant noticed. As such, lights could not simply be left on, and had to be triggered as needed. Second, lights should not turn on or off in-scene, as this change itself is a confounding stimulus (Franconeri and Simons 2005), and should change off-camera. Third, there should be a consistent minimum delay between the stimuli (but not fixed, to avoid predictability), to avoid confusion over which light a participant responds to (we used one second). Finally, light occurrence between the videos in a collage should be evenly balanced.

The coordination of lights turning on and off within a video, and between videos, was non-trivial, particularly given how videos would be combined into various collage configurations. I employed a master schedule that dictated light timing and made minor imperceptible changes to video speed to ensure all constraints were met. Each video was over six minutes and four seconds long and had exactly 12 light stimuli. Thus, each collage had exactly 48 light stimuli, which showed up on average every 8 seconds. As each video had a unique stimuli timing, the relative timing between the stimuli changed in each collage due to my Latin Square balancing.

Cue Integration into Video

All visual cues were created using post-processing in Adobe Premiere and After Effects. For each experiment, a full set of collage videos were made for each cue (each collage had a version with one cue type applied) to allow for within-participant counterbalancing.

Figure 3.3. We compiled five different (but comparable in character) four-tile collages for a repeated-measures study design. We modified the video selection and position in the collage using incomplete Latin Squares to minimize learning effects. Each video and collage lasted six minutes and four seconds long.

Rather than simply attaching cues to all lights in a video collage, for improved ecological validity I also included false-positive cues (cue without stimuli), false-negative cues (stimuli but no cue), and cue misses (a stimulus, but cue at an incorrect location). These not only simulate the realities of imperfect saliency-detection systems (Borji et al. 2015), but were designed to imbue a sense of diligence in participants, as they could not completely trust the cueing system.

Further, introducing cueing errors into a system, while realistic, can have overall detrimental effects on performance: operators can lose trust in unreliable cues and overcompensate with increased attention, introducing additional error (M. I. Posner, Nissen, and Ogden 1978; Michael I. Posner 1980; Abrams and Christ 2006; Franconeri and Simons 2005), potentially more than an un-cued case. As such, these standard errors must be part of a test bed for comparing cues to an un-cued base case.

In my case, each collage contained 52 cue instances: 40 correct true-positive cues (77%), 4 false-positive cues (cue but no green stimulus), 4 false-negatives (green stimulus but no cue), and an additional 4 cue misses, for a total of 23% error cases. False cue rates were based on prior attention work (M. I. Posner, Nissen, and Ogden 1978).

Instruments

Participant taps (indicating they saw a light) produced a short beep to indicate it was registered, and were recorded and automatically processed for response time and accuracy. Accuracy was further broken down into correct identification of a light, tapping with no light or cue, and tapping the cue and not the light in the mis-cue case (cue in wrong location). A tap was correct if it occurred in the correct feed within 2 seconds of the light disappearing (a generous upper limit based on an expected maximum .5s reaction time Michael I. Posner 1980).

After each task (i.e., with one cue), participants filled out a short questionnaire to measure subjective cognitive load (NASA TLX, Hart and Staveland 1988), and custom 20-point Likert-like items (mimicking the appearance of the TLX scales) for nausea, trust in the interface, enjoyability, and self-perception of speed at the task.

At the end of the experiment, participants answered a free-form short answer section on the pros and cons of each cue, as well as any comments on any motion sickness, or other comments.

Participants sat in front of a Microsoft Surface 2 tablet, with the video collage displayed in full-screen and at max brightness, with the minimum tilt setting. Participants were not allowed to pick up the tablet or change the tilt. The desk, chair, and tablet displaying the collages that participants used were placed at fixed initial positions, though participants could adjust the chair to be comfortable.

Procedure

Before beginning the experiment, participants were briefed on the task before reading and signing an informed consent form. They were then given a 30-second practice collage to watch and shown how to indicate where in the collage a light appeared (by tapping). They were told that the videos were pre-recorded using real robot, and were informed of, with examples, of how the cueing system sometimes made mistakes (false cues).

Before starting the tasks, participants are shown example collages containing all visual cues in the order they would appear in the experiment. Participants were told (and reminded before each task) to act as quickly as possible as time was being recorded.

The experiments used within-participants design, with participants completing the task with all cue designs; cue orders and cue-collage mapping were counterbalanced using incomplete Latin Squares.

Before each task I displayed the cue to refresh the participant’s memory, and the task started when the participant touched the screen). Between tasks, before moving on to a new que, there was a mandatory three-minute break to mitigate the impact of fatigue; during this time, participants filled out the post-task questionnaire.

After all tasks were complete, participants completed the post-experiment questionnaire, and were debriefed on the experiment.

3.5 Cue Design

My cue-design methodology was based on my background exploration in human attention literature (summarized in Section 3.1.3), combined with my video game inspiration. As motion is highly effective at drawing attention (Abrams and Christ 2003; Abrams and Christ 2006), it is a strong candidate for cue design. However, motion can be distracting (Tasse, Ankolekar, and Hailpern 2016), and may have a negative impact on primary task performance. Further, in my search and rescue application, the visual field is already noisy: constantly changing as the robot navigates; I need to investigate if motion cueing is still effective in this scenario, or, if the combined motion of the cue and robot becomes even more distracting. I investigate cue motion as a design variable: cues that move (*moving* cues), and cues that do not move (*static* cues). As the light cue is always moving in the visual field of the robots, I defined static cues to be fixed relative to the moving light (and moving visual field).

On-screen cue location is important as it impacts the cue visibility: an operator may be focusing elsewhere when a light appears. Cues located near a light encode the location of the light and thus reduces the search space once noticed (Stainer, Scott-Brown, and Tatler 2013). Therefore, we may expect these cues to elicit fast response times, as once a cue is seen, an operator does not need to search for the light. However, due to inattentional blindness, operators may not notice even highly salient cues outside their current attention (Simons and Chabris 1999). As such, I examine full-screen cues (visible everywhere at once). These should be easy to notice, no matter where an operator is focusing, which indicates to the operator that a target is currently on screen. Therefore, I investigate the cue proximity as my second design variable: cues at the light stimulus (*at-light* cues) and cues over the entire visual field (*full-screen* cues).

I use these two design variables, cue motion and cue proximity, to drive my cue design as well as evaluation.

3.6 Initial Cue Design and Pilot

I conducted an initial pilot study as a broad exploration into cue design for supporting teleoperation visual search, using my two design variables: cue motion, and cue proximity. In the pilot, my full protocol was not followed: I only measured accuracy, and I used an earlier and rougher video collage that was longer, had less rigorous light spacing, and had all cue types intermixed.

3.6.1 Initial Cue Design

I designed and implemented an initial set of nine cues based on my perception literature exploration and my two design variables.

My initial at-light (cue proximity variable) cues were *red circle* and *grey circle*, simple outlines, and *exposure*, a disc of increased exposure (thus brighter image), over the stimulus. These were chosen to explore the impact of visual contrast, a known factor in salience (Howard et al. 2009; Donald, Donald, and Thatcher 2015).

For investigating movement, I animated the grey circle to bounce one cue radius either left to right (*horizontal* cue) or top to bottom (*vertical* cue). Motion direction, either parallel to or orthogonal to visual flow, can impact salience (Franconeri and Simons 2005; Abrams and Christ 2003); given that my robots turn often but do not look up or down frequently (except when they fall), *horizontal* is parallel and *vertical* is orthogonal to expected visual flow.

For the full-screen component of cue proximity, I aimed to impact the whole visual field, to be difficult to ignore, while simultaneously trying to indicate where the light is. I tried blurring (*blur* cue) or darkening (exposure reduction) the entire screen except for a disc around the light. Both changes in clarity and exposure have been shown to be salient (Bridgeman, Hendry, and Stark 1975; Veas et al. 2010). For these cues, I only implemented static versions; simply animating them would not be effective as both the blur and darken effects had no global variation and thus would not show change as they move.

For my full-screen, moving cue, I drew from video-game design and implemented a common targeting animation: *target* was a circle approximately the size of the screen that appeared and rapidly shrank towards the target. While the shrinking motion should attract attention, particularly as a shrinking circle has orthogonal motion to all visual flow directions (Franconeri and Simons 2005; Abrams and Christ 2003), a risk is that it may appear as a receding motion, which has been shown to be less salient (Abrams and Christ 2003).

For all moving cues, the animation lasted for 1 second for consistency across cues; horizontal and vertical bounce cues and target cues all became static grey circles until the light left the screen.

3.7 Pilot Study

My primary focus of the pilot study was to direct my exploration for more formal study by evaluating my test bed, testing my design variables, and obtaining an initial sense of my cue design successes and failures. As such, I ran my pilot with all eight initial cues (red circle, grey circle, exposure, horizontal bounce, vertical bounce, blur, darken, and target) and compared their impact on how many lights participants correctly identified (accuracy). While I compare data across all cues, of particular interest was which cues performed best in each design configurations: at-light static, at-light moving, full-screen static, and full-screen moving.

In addition, I added a case with no cueing (just the light stimuli): this was to measure the overall impact of cueing (e.g., can possibly make performance worse), as well as to test the task itself, to ensure that it was sufficiently difficult to benefit from cues.

3.7.1 Results

I conducted my pilot with nine participants recruited from my general university population. The statistical results were used to inform the next phase of my study, although small sample sizes made my conclusions not as robust as it could be. A one-way repeated measures ANOVA showed an effect of cue type on accuracy ($F_{8,64}$=5.71, η^2=0.42, p<.001, Figure 3.4). Post-hoc comparisons (with Bonferroni correction) comparing all cues to each other revealed the best performing cue for each design parameter set. While Target was the only full-screen moving cue, it was validated by performing statistically better than four other cues. If there was not a clear winner in a design category, I simply picked the cue with the highest average mean accuracy.

In the no-cue case, operators found on average 66% of lights (std. dev 17%). This was comparable to some of the worse cues (exposure, horizontal bounce). This indicates that my test bed and visual-search task are sufficiently difficult, where the addition of cues can potentially improve on the success rate. However, some cues seem to perform at least as badly as having no cue at all (Figure 3.4). In the next section, I detail how these results were integrated into a full formal experiment to better understand the effects of each design.

3.8 First Design Iteration

The results of my pilot study gave me initial representative candidates for each design parameter combination: grey circle performed best as my at-light, static cue. The full-screen dark effect was my best full-screen, static cue. The vertical bouncing circle was my best at-light moving cue. Finally, the video-game inspired targeting was my best full-screen moving cue. I refer to these as *circle, dark*, *bounce*, and *target*, respectfully, as summarized in Table 1 and shown in Figure 3.5.

We employed the full test-bed protocol (Section 3.4.1) as a within-participants experiment: each participant completed the task with each of the four interfaces. I counterbalanced cue and collage order.

Table 3.1. Representative cues for testing our cue proximity and movement design parameters.

		cue proximity	
		at-light	full-screen
cue movement	static	circle	dark
	moving	bounce	target

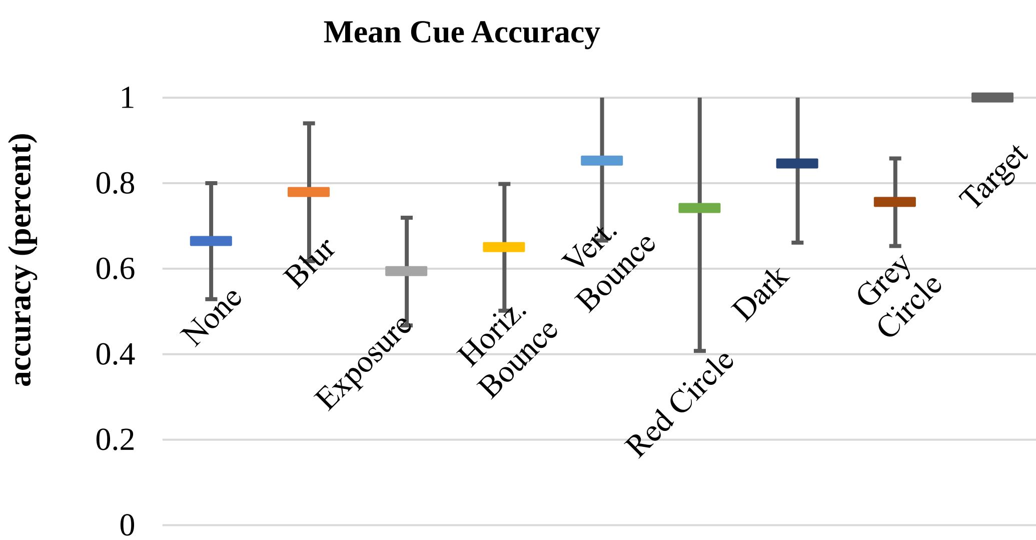

Figure 3.4. Mean accuracy of our initial cues. Error bars show 95% confidence intervals.

The purpose of this iteration was to more formally and rigorously test my design variables, cue proximity and movement, using my candidate representative cues developed through the exploratory pilot study. I again included the no-cue case to more rigorously test the overall impact of cueing in comparison to the un-cued base case (e.g., cueing may possibly hinder performance). Specifically, I have two hypotheses: my cues will outperform a no-cue situation in all measures, which I will test with a repeated measures 1-way ANOVA with planned-contrasts; and that both cue proximity and cue movement will affect operator performance. This latter hypothesis is a 2x2 design, and leveraging my within-participants design I will perform a 2x2 repeated measures ANOVAs on both variables.

3.8.1 First Iteration Results

I recruited 20 participants (8 female) from the local university population. The mode age (collected in ranges) was 26-30, at 35%.

A repeated-measures one-way ANOVA comparing all cue against the no-cue case showed an effect of cue type on response time (Figure 3.6b, $F_{2.8,52.3}$=41.9, η^2=.69, p<.001, Greenhouse-Geisser correction), accuracy (Figure 3.6c, $F_{2.0,38.3}$=30.8, η^2=.62, p<.001, Greenhouse-Geisser correction), and cognitive load (Figure 3.6a, $F_{2.2,41.8}$=6.5, η^2=.26, *p*=.003, Greenhouse-Geisser correction). Planned contrasts against no cue showed all others to be more accurate and to have lower cognitive load (*p*<.002), while circle, bounce, and dark had faster response time (*p*<.01); no response-time difference was found against target (*p*>.05). While Figure 3.6 shows overall means and confidence intervals, the within-participants statistics uses relational scores.

I performed 2-way repeated-measures ANOVAs (cue proximity by motion) on operator accuracy, response time, and cognitive load (Figure 3.7). Bonferroni-corrected post-hoc tests were performed (with main effects) to investigate the effect for each of the levels.

I found a main effect of cue proximity on operator response time: at-light was faster than full-screen (Figure 3.7b, $F_{1,19}$=107.3, η^2=.85, *p*<.001, 95% CI [-232ms, -154ms]). Post-hoc tests revealed that circle was faster than dark (*p*=.021, 95% CI [-153ms, -14ms]), and bounce was faster than target (*p*<.001, 95% CI [-353ms, 252ms]).

I also found a main effect of cue movement on response time revealing that static was faster

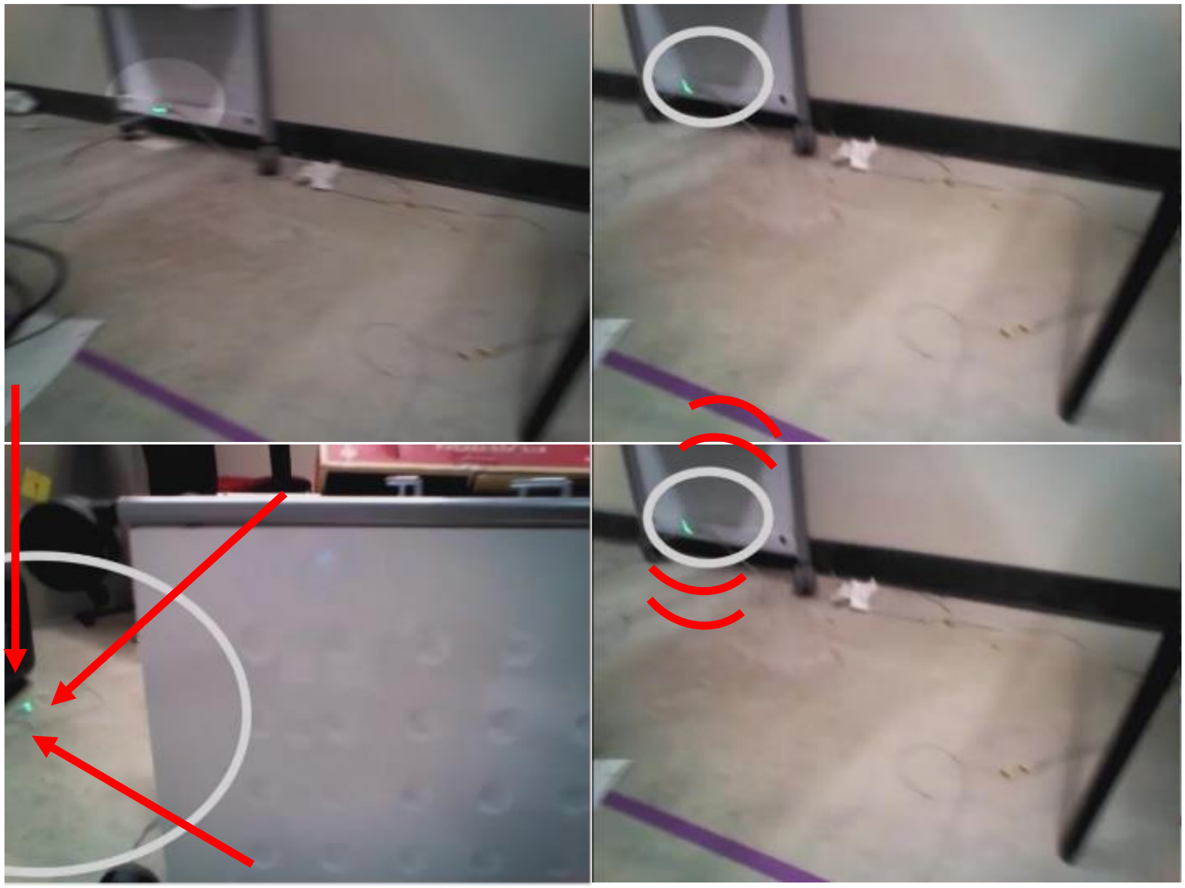

Figure 3.5. Our four interfaces for the first design iteration: dark (top-left), circle (top-right), target (bottom-left), bounce (bottom-right). Red markup indicates motion and are not shown in the interface.

than moving, (Figure 3.7b, $F_{1,19}$=4.9, η^2=.20, p=.04, 95% CI [-113ms, -3ms]). Post-hoc tests revealed that dark was faster than target, p < .002, 95% CI [-235ms, -98ms]; static circle versus moving bounce was not-significant. There was an interaction effect between the two parameters ($F_{1,19}$=24.3, η^2= .56, p<.001).

We found a main effect of cue proximity on operator accuracy, revealing that operators found more lights with at-light cues than with full-screen cues (Figure 3.7c, $F_{1,19}$=4.4, η^2=.19, p<.05, 95% CI [0 lights, 1.15 lights]). Post-hoc tests showed circle to have better accuracy than dark (p=.42, 95% CI, CI [-.23, -2.8 lights found]).

I also found a main effect of cue motion on operator accuracy: operators found more lights with moving cues ($F_{1,19}$=6.5, η^2=.26, p<.05, 95% CI [-1.773, -.177 lights]). Post-hoc tests were all n.s.

I found a main effect of cue proximity on cognitive load, revealing that participants rated full-screen cues as demanding lower cognitive load (Figure 3.7a, p<.01, $F_{1,19}$=8.4, η^2= .31, 95% CI [-9.087, -1.463] NASA TLX points). Tests for main effect for cue movement, and interaction effects were all non-significant.

In summary, the 2x2 ANOVAs indicated that participants were faster and found more lights with at-light cues than full-screen cues, although the full-screen cues demanded lower cognitive load. Participants further were faster with static cues than moving cues, although they found more lights with moving cues.

No effects were found on nausea, trust in the interface, cue enjoyability, preference, or self-perception of task speed. Further, data on miscues and mis-clicks (when no light was present) were all non-significant.

3.8.2 Analysis of Participant Feedback

To gain insight into strengths, weaknesses, and differences between cues, I performed open coding on short-answer feedback. My open coding attached themes to participant feedback, and grouped and iterated on those themes repeatedly to discover core concepts (Berg 1989).

Motion cues (bounce and target) were seen by some as attention-grabbing (positive, 16 participants for bounce, two for target), while distracting by others (11 for bounce, seven for target), describing them as being "tiresome," and "stressful to eye," (bounce) or "breaking visual concentration," and "very distracting" (target). Some participants simultaneously praised the attention-grabbing properties of bounce while commenting that it was too distracting. Static cues had fewer comments on distraction: four participants made comments such as the dark cue being "somewhat distracting," but seven noted that it was "minimally distracting." 11 participants noted that circle was "not easy to locate," and "not very distracting."

Participants commented that full-screen cues positively affected their comfort. Eight participants mentioned that the dark cue was "very relaxing," induced "less dizziness," and was "easy to visualize," and target cue had five comments about reduced stress "it highlights and easy to see. Less effort." Circle had six comments note it was "calm with a smooth motion" and that "appearing without a sense of movement diminished the urgency." No similar comments were given for bounce, but 11 participants complained it was tiring: "made me feel dizzy," "have to cautiously monitor four screens. Stressful" and, "heightened the sense of

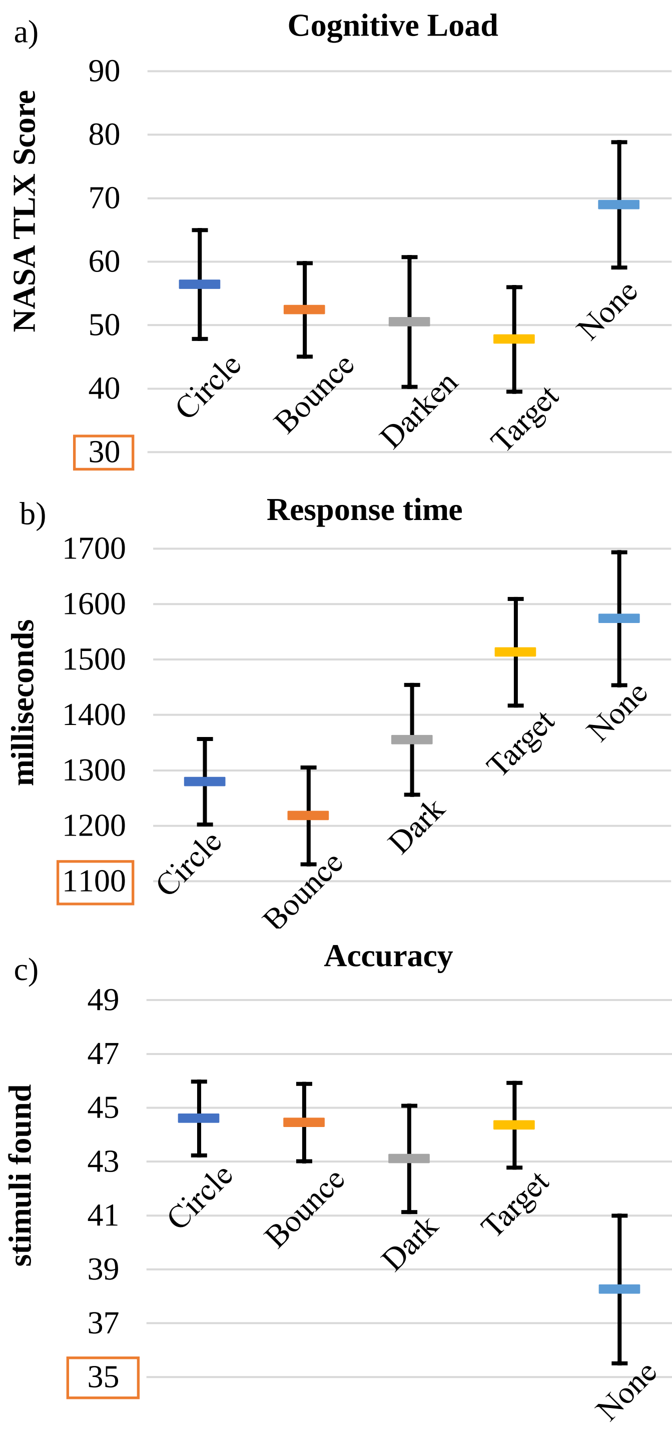

Figure 3.6. Results of planned contrasts, error bars are 95% confidence intervals. a) Cognitive Load Sum (range [5,120]): all cues performed better than no cue (p< .002). b) Response time: Only target was not better than no cues (p< .001). c) Mean Accuracy (range [0,52]): all cues performed better than no cues (p< .001).

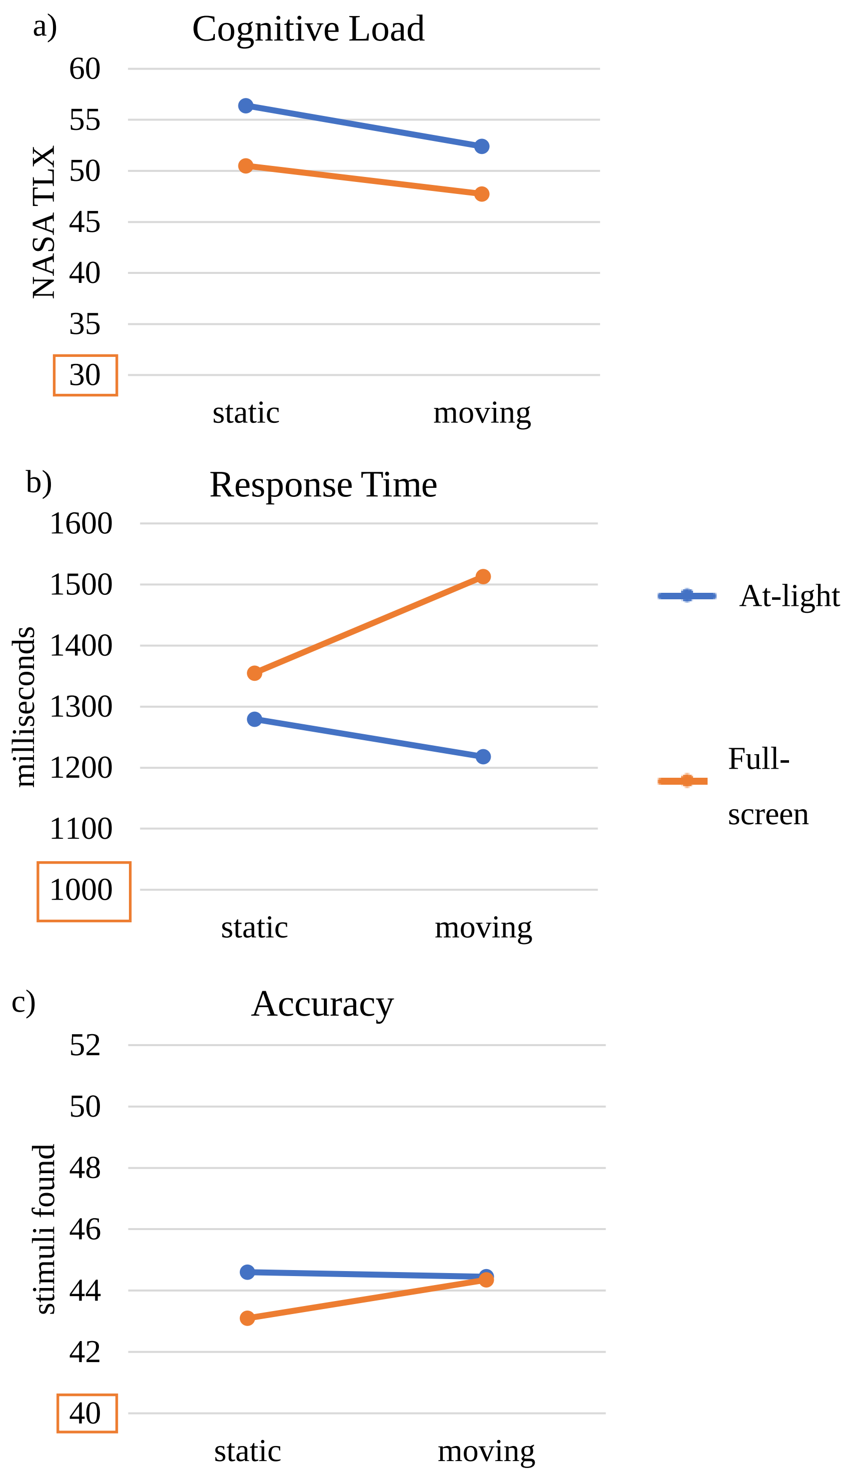

Figure 3.7. 2x2 (cue proximity X movement) results. a) Cognitive load (range [5, 120]): at-light cues had higher cognitive loads (p=.009). b) Response time (in milliseconds): at-light cues were faster (p< .001), and moving cues were faster (p=.04). An interaction effect is clear (p<.001). c) Accuracy (range [0,52]): at-light (p=.05) and moving (p= .019) had a higher accuracy.

urgency."

Unique to the target cue, participants complained about its speed: five participants mentioned frustration: it was "too slow to capture the light," and "took focus off other screens for extended periods." No other cue had comments about speed.

Some participants noted the light-position information encoded in the full-screen cues. 10 mentioned that the target cue "helps you to target your focus on one [robot's video]" and that it "can tell me directly where the green light is when it's shrinking." While four people mentioned that with the dark cue the "increase light-to-background contrast made it easier to detect," eight participants conversely mentioned that it only "darkens the whole screen to let you know something has been captured" and required participants to "waste time by locating the image that is less darkened." There were no such comments found for the at-light cues.

Related to comfort, 14 participants complained about some level of nausea from the study. This was not linked to a specific interface, but was attributed by the participants to the simultaneous and often not coordinated movements of the robots.

3.8.3 Insights on our Designs and Design Parameters

As with the pilot, my results confirm the validity of my test bed, as well as my cueing technique: when compared with no cue, all cues increased accuracy and lowered cognitive load, and all but the target cue improved response time (though target was not found to be worse). Therefore, at the least, cueing may be useful to help participants in urban search and rescue tasks. Further, the benefits of my cues, for the results I measured, offset any detrimental effects that may be introduced, such as too much distraction (Tasse, Ankolekar, and Hailpern 2016).

The analysis of my two design parameters detailed important tradeoffs between design choices. Full-screen cue proximity appeared to demand lower cognitive load than at-light cues. It is not entirely clear why this may be, but participant feedback indicates that some found full screen cues more comfortable and readable, with target specifically being helpful for directing attention to a light. As well, bounce's motion (quick bounces up and down) is different from target's (smooth shrinking) which may also be a factor. Further, full-screen cues may reduce stress of potentially missing a cue, as they are much more difficult to miss.

At-light cue proximity resulted in operators finding more lights, and more quickly, than full-screen cues, despite the increased cognitive load. This can be explained in part by how at-light cues effectively make the stimulus larger. Further, at-light cues immediately indicate where the light is (once the cue is noticed). This can be contrasted with dark where participants complained they had to take time to search for it, or target, where they complained of the slow speed of my moving full-screen target cue which took an entire second to home in on the light.

Note that moving cues had slower response time on average. This is supported by my planned contrasts on response time: all interfaces except for target (a moving cue) were faster than None. Further, Figure 3.6b and Figure 3.7b suggests target is the driving force behind the response times, performing worse than other cues, except for the no cue case.

For the cue motion design parameter, participants found more lights with moving cues in general, and overall were finding the lights faster than with static cues, although the specific results were mixed (the static dark was faster than the moving target). Participant feedback indicated that the moving cues were more salient, which explains the improved accuracy. The attitudes were again mixed, however, with some framing this positively as attention-grabbing or negatively as distracting, although the negative component was not reflected in performance or cognitive load scores.

Static cues were in general the poorest performers, with higher cognitive demands, and slower response times, and some participants noting that they were not easy to locate. This is supported by dark being specifically worse in accuracy. Participants commented that they could easily see the dark cue, but had to quickly look for the light itself: there was no position information encoded in the cue. This can be compared to target which directly led the viewer to the light by the end of the animation.

Moving towards a Second Design Iteration

While the circle cue was better than no cue, and helped operators find the most lights, it appears to have one of the highest cognitive load demands, and did not perform well in response time (except or being faster than target) or accuracy. The dark cue similarly only performed well for speed in comparison to target, and had the lowest accuracy. Bounce was one of the strongest performers, with one of the fastest response times and best accuracies, but had complaints about distraction and fatigue, with improvements to be made in cognitive load. Target, while having speed issues, had one of the stronger cognitive load scores and was comparable on accuracy.

One important component of my analysis was the interaction effect found on response time (Figure 3.7b), which appears to be driven by the slow target cue. Target’s poor speed was also reflected in it being the only cue that did not perform faster than no cue. Thus, I need to be careful about interpreting the response time main effects, as the specific target cue may need design improvements.

3.9 Second Design Iteration

My main goal in the second iteration was to develop hybrid cues with both static and moving elements, as well as full-screen and at-light elements to see if combining my cue design parameters could improve operator performance. I draw on my first design iteration results to develop a new set of three hybrid cues to be tested. Some of these are iterations on my previous designs, while others are new designs based on my results from the first iteration.

While the bounce cue was a strong performer, the weak point was the fatigue and cognitive load. I iterate on bounce by adding a full-screen element to try and mitigate these issues, to embed the comfort and cognitive load gains associated with my other full-screen cues. Specifically, I add a border to the video feeds for the duration of the cue, without changing the bounce itself (Figure 3.8). I hypothesize that this *framed bounce* will have improved cognitive load scores over the previous bounce, without negatively impacting response times or accuracy.

Participants commented on the benefits of full-screen cues encoding the location of the stimulus as well as providing an alert. I designed a new full-screen cue that statically encodes the light position; I hypothesize that avoiding moving elements can reduce frustration (and help with cognitive load), while maintaining the accuracy and cognitive load benefits of the full-screen design. Specifically, I used a greyscale radial gradient (linear, dark at edges, light at center) centered over a light (Figure 3.8). I anticipate that eyes can quickly follow the

Figure 3.8. Our two new cues (from the left) framed bounce, and tunnel. Red lines indicate the animation

gradient from anywhere on screen toward the light as if looking through a tunnel. This encodes location similar to the target cue, but without the time constraint. I hypothesize that this *tunnel* cue will improve the cognitive load and accuracy over bounce, while achieving similar with response time.

As a secondary agenda, I iterate on my target cue. I believe that the animation that target uses to shrink toward the stimulus could be much faster, and still maintain its positive characteristics (low cognitive load and strong accuracy). As such, I developed a *fast target* variant which animated three times faster (0.33s instead of 1s). I hypothesize that this improved target cue will maintain the accuracy and cognitive load of target (not do worse), while improving on the response time.

I conducted my full protocol using five cues: framed bounce, fast target, tunnel, regular target (to compare against fast target), and regular bounce (to compare against framed bounce, and tunnel). I keep target and bounce in the procedure to re-test for consistency and improved comparability with within-participants. Due to testing specific designs against prior iterations (comparing target to fast target) or best performers (comparing bounce to framed bounce and tunnel), I used *t*-tests to compare methods for each measure.

3.9.1 Results and Discussion

I recruited 20 people from around my university campus. Participant ages were collected in ranges; the mode was the 18-20 range, at 55%. I analyzed my data using *t*-tests given my targeted hypotheses and did not use more exploratory methods. I summarize these results in Figure 3.9.

When comparing framed bounce to bounce, I found a trend for framed bounce to improve cognitive load (t=1.6, p=.064, 95% CI [-1.582, 11.682], one tailed). While no difference was detected between response times (t=-1.7,p=.11,[-96ms, 10ms]), framed bounce had better accuracy (t=-2.5, p=.021, 95% CI [-2.290, -.210 stimuli]).

Comparing tunnel to bounce, I found a trend for tunnel to improve cognitive load (t=1.5, p=.08, 95% CI [-1.771, 9.971], one tailed), with no difference found on accuracy (t=1.0, p=.16, 95% CI [-.725, 2.125 stimuli], one tailed). I found tunnel to be slower than bounce (t=-3.6, p=.001, 95% CI [-162ms, -43ms]).

I found that participants had a faster response time with my fast target, than regular target (t=7.6, p<.001, 95% CI [205ms 361ms], one tailed). I did not find difference for accuracy (t=0, p=1.0) or cognitive load (t=-0.9, p=.19, 95% CI [-8.503, 3.303]).

In this study, I successfully demonstrated how hybrid cues can be developed to integrate benefits of cues throughout my design space. My full-screen plus at-light framed bounce cue had better accuracy than the regular bounce, comparable response times, and potentially improved cognitive load (a trend, requiring further study). At the same time, the failure of the tunnel cue, which may slightly improve cognitive load but harms response time, highlights the non-trivial nature of designing effective cues. It is likely my position encoding failed somehow, as it is a core concept in target, which performs well. Finally, I have demonstrated how my target cue can be improved simply by making it faster, negating many of the problems encountered with this cue in earlier study though my data points to a potential small effect of increased cognitive load.

3.10 Reflection on our Design Process

Looking across both studies (Table 3.2), motion and full-screen cues seem to be effective at improving accuracy, response time, and cognitive load. I also made at-light cues with good accuracy and response time, implying that, at least in my experiment scenario, people are good at searching for local cues as long as they stand out in some way e.g. animation. In both studies, there were moving (original target) or full-screen cues (dark, tunnel) that did not perform well, which hints at the complexity of the design space; I cannot blindly trust a single or pair of my

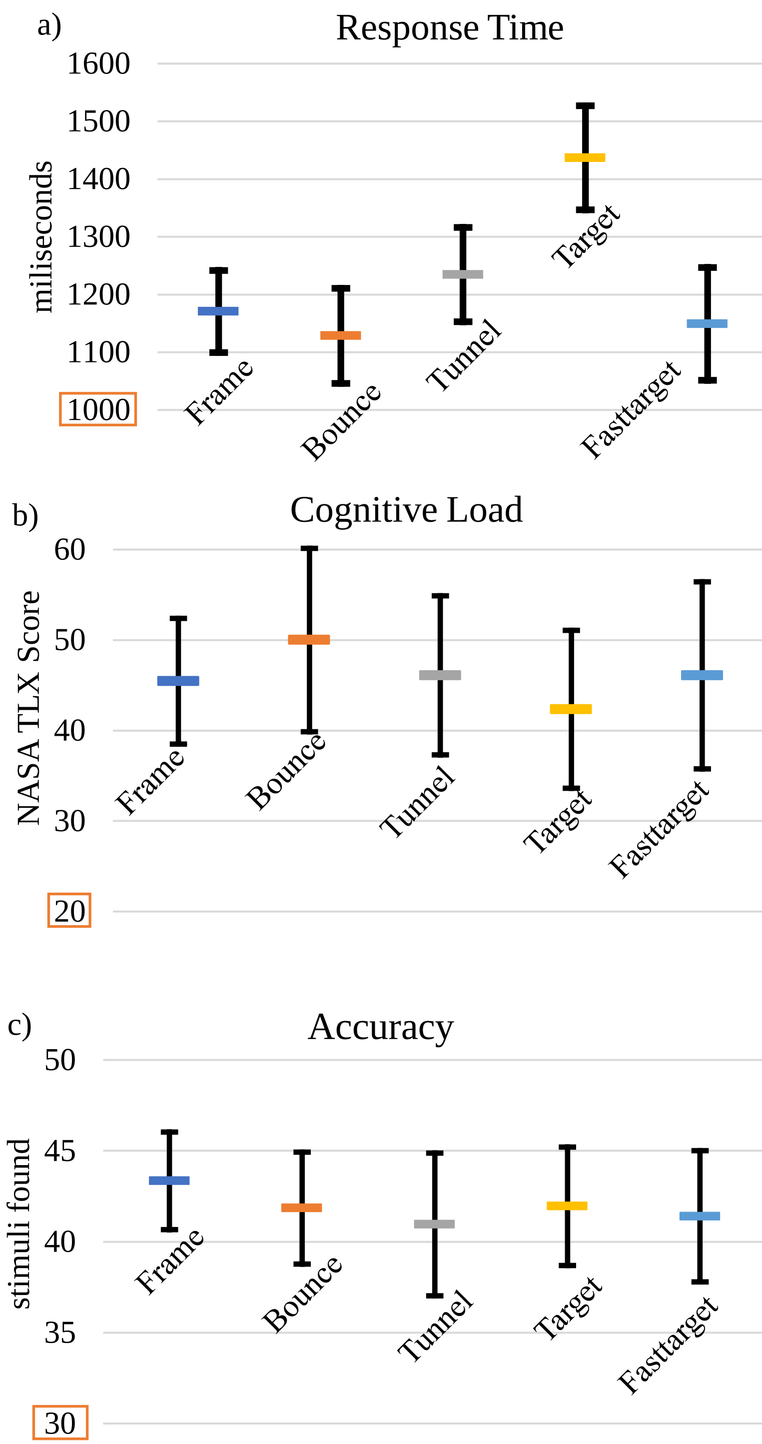

Figure 3.9. T-test results, error bars are 95% confidence intervals. a) Response time: Fast target was faster than target, bounce was faster than tunnel (p< .001). b) Cognitive Load (range [5,120]): there were trends for framed bounce (p=.064) and tunnel (p=.08) to incur less load than bounce. c) Mean Accuracy (range [0,48]): framed bounce helped more than bounce (p< .021).

design variables. I did not see any cue perform worse than no cue, implying that adding visual cues to teleoperation poses little risk, and can add many benefits.

Our process, choosing design parameters and iteratively exploring implementations through performance and user feedback helped me design my best visual cues (fast target, and framed bounce). This is to say, iterative design processes applied to two parameters yielded useful results. Further, studying multiple parameters at once (our 2x2 design) also revealed interaction effects due to design parameters I did not explicitly study (speed). This leads me to suggest that much richer results can be achieved in a complicated scenario, such as teleoperation, by studying multiple variables, despite the challenges associated with conducting such studies.

My work used a scenario that was designed to mimic to a complicated, real-world task. While the lights my scenario could even be considered more obvious than other urban search and rescue search targets, the no-cue case (showing lights without any indicators) proved difficult for operators, and leveraging the perception literature was still effective at improving how often operators found these seemingly obvious lights. The experiment procedure was also a success as the video monitoring method removed the experimental complications involved with controlling multiple robots, while enabling me to easily swap different scenarios (videos) or trying new cue types. I stress, however, my belief that using real robots to record the video is an important ecological consideration. While my study was still "in-lab," I believe my results

Table 3.2. Summary of design parameter effects.

Cue Property	Benefits	Drawbacks
Motion	Higher accuracy	Distracting
Static	-	Poor accuracy, response time, cognitive load
Full-screen	Lower cognitive load	Careful design required for good accuracy and response time
At-light	High accuracy, response time	Increased cognitive load

further validates the attention literature in "messier" ecological scenarios.

False cues were a core component of my scenario design, mimicking previous work, as well as more realistically portraying how real cues in search and rescue would work. A large risk I anticipated was that false cues would undermine trust in the interfaces, and affect performance. Nowhere, however, did I see the false cue rate impact performance enough to counteract the benefit from cueing. Further, no-cue accuracies were low (66-80%) even in my controlled scenario with spaced-out stimuli. This should embolden robot designers to use modern computer vision algorithms to augment their interfaces even if they are moderately unreliable. This technology can improve robotic teleoperation *right now*.

3.11 Limitations and Future Work

While I demonstrated that changing my two design parameters could affect operator performance, I did not explore the full continuum of these dimensions. For example, my motion cues were only animated for a short time. Similarly, target was a full-screen cue, but after converging on the light, it was an at-light cue. It may be that different positions on the dimensions of cue-proximity and motion, may have different results and complex interactions. Exploring these and other design cues in the context of urban search and rescue, such as color, animation speed, length, and even non-visual cues such as sound, will help future interfaces for teleoperation.

I introduced miscues into my data, but no statistical results were found. I believe this is due to low numbers (20% of my stimulI was miscues). In response-critical situations where operators must correctly identify regions of interest that need additional resources (e.g. life-saving medical personnel), missing or mistakenly seeing a stimulus may incur great cost. While I focused on performance measures, future research could target these miscues specifically by longer experiments or greater participant numbers that enable more statistical power, extending the existing visual attention research on miscues to teleoperation.

Another theme that emerged was that cues that grabbed attention could be perceived as either helpful or distracting, similar to previous work (Tasse, Ankolekar, and Hailpern 2016). My experiment hinted that the positive or negative association may be linked to cognitive load –

not task performance – and how a cue could be ignored after it was noticed. This is an important balance to achieve with multiple stimuli on-screen at once (a more realistic scenario than ours), as a cue may be so distracting it distracts from other cued stimuli. Investigating performance with multiple concurrent stimuli and stimuli densities would better illuminate how cues can draw attention away from multiple video feeds.

Finally, this research was in the context of teleoperation, but no robots were actually controlled. Single robot teleoperation remains a challenging open problem that will take much of an operator's cognitive resources, so exploring visual attention while actually controlling a robot is important future work. Moving away from teleoperation, my work may be further generalized by comparing visual search with a still camera and moving target as opposed to my moving camera with still targets.

3.12 Conclusion

In this chapter, I introduced my investigation of video game inspired attention drawing cues in a multi-robot control context. Specifically, I borrowed design ideas such as full-screen damage indicators, smaller targeting indicators, and animations, and these designs also leveraged and agreed with the existing literature on attention.

To explore my video game inspired designs and two design variables of cue proximity and cue motion, I designed and evaluated seven different visual cues through an iterative design process. In my mock-disaster scenario, I found my search task was sufficiently difficult to warrant study, and my cues proved useful (improving accuracy, reducing workload, etc). My design parameters had tradeoffs in performance and cognitive load, and my results indicated that full-screen and animated cues can improve accuracy, response time, and cognitive load if they are designed well. My research provides a baseline for more research to understand cueing visual attention in teleoperation.

Overall, my results demonstrate, as a proof of concept, how video game interface designs can be leveraged to improve teleoperation task performance as well as operator experience. My results additionally highlight how video game techniques can align with knowledge in the psychology literature; in this case, I focused on how game interfaces could aid attention by reducing the cognitive load of visual search.

Video games help players reduce cognitive load in multiple ways; they provide players visualizations to stay aware of their character's state, help players understand and navigate complex environments, or enable players to better understand and quickly react to their surroundings (see Chapter Six). Thus, my chapter acts as an example of how video game interfaces like these can be adapted to teleoperation, and how they may benefit operators from a variety of video game interaction techniques.

CHAPTER FOUR: SHAPING TELEOPERATOR DRIVING BEHAVIOUR AND PERCEPTIONS BY PRIMING

Video games are inherently virtual, and must convey a sense of physicality of the objects and characters within it. For example, in racing games, cars will accelerate differently, may drift differently in turns, or maybe slip differently on pavement because of their virtual weight. Games encourage their players to feel or believe these physical properties, even though they are not real – games have even tricked players into thinking there were differences when two virtual vehicles act identically! For example, in Hi-Octane [Bullfrog, 1995], each vehicle has varying statistics presented to the player such as top speed and acceleration, but the developers admitted that all vehicles drive the same – despite this, players believed the displayed driving ability (Scheurle 2017). I take inspiration from how games shape player perceptions of physicality in a virtual world and explore how interaction design can shape an operator's perception of the physical abilities of a robot, and how they operate it.

In this work, I explore shaping people's teleoperation behavior and expectations about robot capability with *priming*. I introduce the concept of priming, borrowed from psychology, and present two novel priming techniques for teleoperation. In a series of four experiments, I investigate my priming techniques and the interactions between them. My results found that priming can affect an operator's perception of the robot including its speed, weight, or overall safety, and in some cases can encourage operators to drive more safely. Given these results, I argue that interface designers should consider how their robot design and its presentation (marketing materials, instruction booklets, etc.) may prime operators to shape expectations about robot capabilities and impact how they teleoperate a robot.

My results show how even non-functional design changes – those that do not impact how the robot behaves or responds – can impact the operator's behaviour and experience. By borrowing aesthetic presentation techniques from video games, I can influence how people engage with teleoperation, changing how safely they drive, how much workload they experience, and how they perceive the robot they control. Thus, video games can also teach us how to change perception and expectations in teleoperation.

Parts of this chapter have been taken from the following publications:

Daniel J. Rea, James E. Young, "Methods and Effects of Priming a Teloperator's Perception of Robot Capabilities." The ACM/IEEE International Human-Robot Interaction Pioneers Workshop. ACM. 2019. (31% acceptance rate)

Daniel J. Rea, James E. Young. "It's All in Your Head: using priming to shape an operator's perceptions and behavior during teleoperation." The ACM/IEEE International Conference on Human-Robot Interaction (HRI'18). ACM. 2018. (23% acceptance rate)

Daniel J. Rea, Mahdi Rahmani, Neil Bruce, James E. Young. "Tortoise and the Hare Robot: Slow and steady almost wins the race, but finishes more safely." IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN '17), IEEE. 2017.

4.1 Driving and Expectations in Teleoperation

One common problem in teleoperation applications is navigation – operators driving robots must be careful of injured people or damaged buildings in urban search and rescue, sensitive equipment in industrial inspection tasks, or everyday people during telepresence navigation. Navigating safely is, however, challenging, and research suggests that operation mistakes and errors, or *critical incidents* such as collisions with people or objects, are commonly due to operator error (Giese, Carr, and Chahl 2013; Williams 2004). Thus, a continuing body of work explores how to design teleoperation interfaces in order to support operators and encourage safe and effective teleoperation.

In research on controlling motor vehicles, changing how the vehicle is presented to the driver has been shown to impact perceptions of the vehicle's abilities, safety, and change how people drive. Intuitively, we see this when people may drive more safely if they expect dangerous situations (Fuller 2005) such as icy roads, or if we think the brakes to be in poor repair or the vehicle may stall; we drive differently even if none of these conditions are actually true. Inversely, a person may drive a car with more safety features less carefully (e.g., as with ABS brakes in the past, Jonah, Thiessen, and Au-Yeung 2001). I explore if the perceived driving qualities of a teleoperated robot can similarly influence how an operator drives the robot and how they perceive its abilities and safety.

Video games use this idea as an explicit design goal, leveraging perception or expectations in their design to affect player perceptions and behaviour. Games influence perceptions by hinting at, or outright describing physical qualities, such as with specification screens (Figure 4.2, Figure 4.1), which can influence how a player approaches the game. For example, in Figure 4.2, the go-kart appears to have slow acceleration, high top speed, and high weight. This gives the suggestion of a tough vehicle with a large and powerful motor. This further informs the user to drive carefully, as losses of speed may cost them time (due to slow acceleration), encouraging certain styles of driving.

However, not all changes in the presentation characteristics make changes to how the kart drives, so while the player believes an upgrade is making the vehicle better, it is not necessarily so. In the game, increasing a certain driving statistic such as acceleration will increase the visual representation of that number, but not necessarily change how the go-kart physically drives ("Mario Kart 7 In-Game Statistics" 2019). This is taken to the extreme in the game Hi-Octane [Bullfrog, 1995, Figure 4.1]. In this game, each vehicle has different physical abilities displayed, such as speed or armour, but the developers have revealed that all vehicles drove exactly the same, which appeared to have gone unnoticed even by avid players (Scheurle 2017). Video games take this idea to a further level, orchestrating entire experiences in a game by targeting player expectations. Ominous music may play to warn a player of an incoming boss fight [e.g., Mass Effect 2, BioWare, 2010] or to scare players [e.g., Resident Evil 4, Capcom, 2005; Bioshock, 2K Games, 2007]. Further, video games may encourage aggressive playing with fast-paced metal music [e.g., DmC: Devel May Cry, Ninja Theory, 2013], or

Figure 4.2. A go-kart specification sheet in Mario Kart 7 (Nintendo, 2011). The stats shown on the left suggest physical qualities of the vehicles, in this case, a slow, heavy, but powerful vehicle. In reality, these displayed statistics do not always accurately map to vehicle performance, but are designed to give the impression of a vehicles driving ability or a sense of progression.

Figure 4.1. A vehicle specification sheet in Hi-Octane (Bullfrog, 1995). Vehicles had wildly varying stats, but the developers have revealed that the stats did not reflect vehicle performance, and all vehicles acted exactly the same.

convey power by creating a large controller rumble when firing weapon [e.g., Destiny 2, Bungie, 2017]. I aim to leverage this form of presentation to encourage operators to believe particular characteristics of a robot, even if they are not true, and see how that changes their perception and driving of the robot.

In human-robot interaction, we know that a person's expectations and perceptions of a robot and its abilities can heavily impact aspects of the user experience such as shaping user expectations, interaction satisfaction, and motivation, as well as being able to influence the adoption and acceptance of technology (J. E. Young et al. 2009; Lindgaard et al. 2006; Mitra and Golder 2006). Research in human-computer interaction has found that such primed perceptual changes may change behaviour (e.g., Banakou, Groten, and Slater 2013; Thellman and Ziemke 2017). Even without performance differences, priming can be an important tool and concept for roboticists to shape and control how their robots are perceived and accepted.

I posit that explicitly and strategically designing for teleoperation experience to seem purposely more or less safe can shape operator psychology (e.g., frustration, confidence), opinions of the robot, operation satisfaction, and even their driving behaviour (Figure 4.3). Specifically, I investigate *priming*, in the sense of employing stimuli that encourages people to recall past experiences which then influences their thoughts and behaviour (Bargh, Chen, and Burrows 1996; Doyen et al. 2012; Dijksterhuis and Bargh 2001).

We explore three different priming methods, and investigate if I can influence operator expectations of a robot to result in operators perceiving the robot differently and even driving it differently as well. My methods also aim to change an operator's expectations of a robot's capabilities without requiring physical changes to the robot (e.g., capability, durability). Thus,

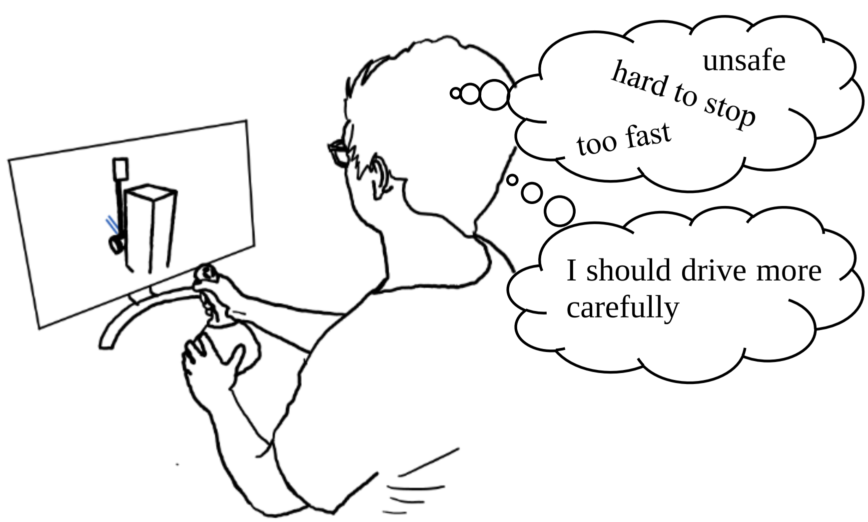

Figure 4.3. We investigate how priming an operator's expectations of robot capabilities impacts their driving behavior and perceptions of the robot. We found changes in perception of the robot in all cases, and, in certain cases, a reduction in collisions.

the goal of this work is to search for low-cost tools that designers can use to improve teleoperation interaction design in addition to improving robot performance.

4.1.1 Priming Design and Results Overview

I present and explore four methods for priming teleoperators to shape their expectations about a robot which I investigated over the course of four experiments (86 participants). Due to the importance of perception of a technology, I evaluate both the impact of my designs on task performance and operator perceptions of a robot.

My initial technique is called *driving profiles*, where I alter a robot's acceleration and maximum speed such that the robot *feels* heavier or lighter and encourages different operator driving behaviour; this mimics the strategies used in video games, where all speed and acceleration curves are virtual and carefully defined by a designer. I designed this to prime people by encouraging them to recall, larger or more dangerous vehicles that accelerate slowly. I found participants drove more safely with a slower robot, and completed the task at a comparable speed to the faster robots. These speed and acceleration changes further changed an operator's perception of their own speed and performance (we published the results in Rea, Hanzaki, et al. 2017). It was unclear, however, if this improvement was due to operators changing behaviour because the driving profile priming changed their perceptions of the robot, or because of the difference in actual robot ability. Thus, a goal of my follow-up studies was to investigate changing perceptions and behaviour without affecting robot capabilities.

My second technique is *tangible priming*, which provides a constant haptic stimulus to the user to represent robot capabilities similar to my driving profile method: in my case, using a stiffer or looser joystick feel (its spring) to represent a lighter or heavier robot. This draws from video games that use haptic feedback to provide a sense of weight and impact in game [e.g., MechWarrior 4: Mercenaries, Microsoft, 2002]. However, in my experiment the spring stiffness in the joystick does not actually change robot ability – I investigate whether the priming alone can encourage different perceptions and operator behaviour, similar to some racing video games (Figure 4.2, Figure 4.1). I found that tangible priming can improve operator driving safety, reduce their perceived workload, and changed an operator's perception of the robot's speed, overall safety, and more (we published these findings in Rea and Young 2018).

The results of tangible priming and driving profiles together lead me to wonder I could produce these effects without a physical component, leading me to my next method.

My next priming method was *descriptive priming*, where I modify my description of the robot to the operator; this draws directly from the descriptive strategies used in video games used to shape player expectations (Figure 4.2, Figure 4.1). I prime by telling participants that a robot is faster, slower, more or less safe, etc. Similar to prior methods, this was to encourage operators to draw from experiences of driving vehicles that share these properties (such as a faster car). I did not actually change robot abilities, but used my priming method to encourage operators to change their driving behaviour because they believe a robot had certain physical capabilities. My descriptive priming was able to make operators believe that a robot was faster, more safe, heavier, better at staying upright, and more, just from how I described the robot to the operator (we published these results in Rea and Young 2018).

Reflecting on my studies and my results, I noted a potential usability confound in my tangible priming case: perhaps a stiffer joystick (our tangible priming manipulation) may simply be easier to control from a usability perspective, explaining the observed safety improvement that was not seen in my descriptive priming results. To follow up on this point, I conducted an additional experiment to investigate whether it was priming (shaping operator expectations) or usability (ease of joystick control) that resulted in safer driving behaviour. To do this, I attempted to remove the priming component by explaining that I was simply changing joystick stiffness without changing the robot (in contrast to the deception used in the earlier studies). My results support my earlier findings: priming (not joystick-stiffness usability) can improve teleoperation performance. However, despite telling operators that their robots did not change, participants still reported that they felt the robots performed differently (we published these results in Rea and Young 2019); this was unexpected. Upon reflection this suggests a much subtler, and more nuanced, role of priming on interaction than I initially expected. I finish my paper with a deeper discussion and reflective analysis on priming, and its use in interface design, stemming from my three experiments.

In short, drawing from my results I propose the intentional use of priming as a new tool for use in teleoperation design. Designers can use priming to help shape operator opinions about their robot and success, and in some cases operator driving safety. I borrowed this approach of influencing expectations through presentation from various video games, suggesting that further non-functional game inspired techniques that affect how people think may be useful in design, such as designing for aesthetics, feel, or fun.

In the remainder of the paper I explore priming background from psychology relating to my work, detail my four studies, and finish with an overview analysis of my technique and results. I briefly outline each priming method and the results of my experiments, and summarize the purpose of each in Table 4.1.

4.2 Background: Priming

In psychology, the term *priming* if often used across a range of applications and methods that uses a stimulus (the priming) to cause an impact on an event or interaction. In my work, I focus on *behavioural priming*, where exposing a person to a stimulus or concept elicits some associated knowledge from previous experience, and impacts their behaviour based on that experience (Bargh, Chen, and Burrows 1996; Doyen et al. 2012; Dijksterhuis and Bargh 2001). For example, showing people a picture of a library can make them unconsciously speak more quietly (Aarts and Dijksterhuis 2003). In this case, people associate the stimulus (an image of

Table 4.1: a summary of our four priming methods and results

Name	Priming Method	Purpose
Driving Profile	Limit robot max speed, acceleration ability	Make the robot feel heavier or lighter by how it accelerates
Tangible	Varying joystick spring stiffness levels	Make the robot feel heavier or lighter by how heavy the control hardware feels
Descriptive	Visual, verbal descriptions of robot capabilities	Make the robot feel more or less dangerous by describing the robot differently
No Priming	Tangible, but no deception	Explore how heavier or lighter control hardware affects usability.

a library) with their prior experience of libraries requiring quiet and change their behaviour to align with that experience (e.g., speaking more softly).

Priming is broadly studied in domains outside of psychology. For example, in marketing it has been shown that priming stimuli embedded in surroundings can change evaluations of a company's brand (Yi 1990), and priming stimuli combined with different prior knowledge was found to change price evaluations of products (Herr 2012). Researchers have also studied potential biological underpinnings of priming in order to better understand the human brain (Iacoboni 2009). How priming may work, and its interactions with other variables is still under study. However, it is clear that priming has the potential to change perceptions and behaviour. Thus, I examine priming as a means of shaping teleoperation.

4.2.1 Priming Methods in the Literature

A broad range of priming methods have been shown to be effective in altering behaviour. Various modalities have been explored such as sound (e.g., Sollberger, Rebe, and Eckstein 2003), tangible methods (e.g., Ackerman, Nocera, and Bargh 2010), or visual stimuli (e.g., Aarts and Dijksterhuis 2003). Explorations of priming effects have investigated the range in stimuli subtlety or frequency (e.g., MacLeod 1989; Forster and Davis 1984).

Researchers have experimented with whether priming needs be subtle or hidden (or can be explicit) to be effective. One famous subtle example showed participants sentences with types of words either omitted or included: in the results, omitting words related to aging induced people to walk more slowly, supposedly because participants focused on those omitted words (Bargh, Chen, and Burrows 1996). Even having a picture in an environment can encourage the associated behaviour (e.g., showing a picture of a library to make people quieter, Aarts and Dijksterhuis 2003; Dijksterhuis and Bargh 2001). Similar priming can be effective even when more explicit, such as playing musical chords designed for specific emotional impact during a task (Sollberger, Rebe, and Eckstein 2003). Priming can even be effective when the person is aware of the stimuli and the expected priming result (Doyen et al. 2012; Cheesman and Merikle 1984).

Priming has commonly been studied in the context of impacting social relations between two people. For example, having one person describe another as “mean” or “kind” can increase the likelihood that the primed characteristics will be observed (Kelley 1950). Opinions of others can also be primed using physical props, such as seeing someone as more important when they are holding a heavier clipboard (Ackerman, Nocera, and Bargh 2010). Priming can be quite nuanced, for example, in the above example people assume the other is more difficult to interact with if the clipboard is rough (Ackerman, Nocera, and Bargh 2010).

Priming stimuli can be presented frequently or continuously throughout interactions (Bargh, Chen, and Burrows 1996; Ackerman, Nocera, and Bargh 2010), and are not limited to being shown only beforehand. For example, driving a car with a loud engine provides ongoing priming, a constant reminder, of the car's power. In the earlier example of the library picture influencing people to speak more quietly, the picture was present throughout the experiment (Aarts and Dijksterhuis 2003). Considering when and how often priming methods are applied should be considered in intentional priming designs.

This breadth of techniques highlights the range of potential priming methods and how they could be integrated into a variety of scenarios. As such, priming should be considered as a component of interaction design.

4.2.2 Priming Effects

Priming is often studied for its immediate effects, but priming can have long-term results, sometimes lasting for hours, weeks, or months (Sloman et al. 1988; Becker et al. 1997). With time, however, the strength of priming effects tend to weaken (Sloman et al. 1988). Repetition of the priming can reinforce the weakening effect, but may not work for all stimuli (Forster and Davis 1984).

Priming effects can be highly context sensitive (e.g., Bargh, Chen, and Burrows 1996; Doyen et al. 2012; MacLeod 1989; Aarts and Dijksterhuis 2003). Indeed, the context or environment itself can be intentionally designed to prime (Aarts and Dijksterhuis 2003; Ackerman, Nocera, and Bargh 2010). For example, priming effects can vary due to the environment (e.g., sounds, Sollberger, Rebe, and Eckstein 2003), or nuances of the task description (MacLeod 1989). Even the experimenter's expectations of the effects of priming may bias or impact results

(Doyen et al. 2012). Thus, measuring the causes and effects priming can be difficult (Westlund et al. 2016; Doyen et al. 2012); how to best study priming effects remains an open question.

4.2.3 Applying Priming to Teleoperation

Priming, in the context of my work, is the use of stimuli that evokes feelings or memories that can affect a person's thoughts or behaviours. The stimuli can be given *before* or *during* a task, may be continuous, or may be done in secret or be explicit. Importantly, priming can even be unintentional, for example by creating the design of an interface that accidentally primes an operator to be more aggressive and dangerous. Thus, I argue it is important for the field to understand priming so as to better enable teleoperation designers to take control of the user experience and make informed decisions regarding how their design choices may be affecting an operator's perceptions and actions.

Although the effectiveness of priming is established in some domains, the science is still unclear on the limits and applications of priming (Doyen et al. 2012). I build on the work of behavioural priming in psychology by extending it to explore the use of priming for shaping teleoperator perception of their robot, and their behaviour. As until now it has been unclear whether priming in general could work in teleoperation, I investigate whether priming can be successful in impacting operator perception of a robot, and whether priming can impact teleoperation performance.

4.3 Priming and Shaping Perceptions to Improve Interaction

Priming and the goal of shaping a user's perceptions and behaviour has been leveraged and studied in a variety of fields. These include traffic psychology, human-computer interaction, and social human-robot interaction. Their findings overall provide evidence that perceived qualities of an object or experience is linked to user behaviour, and that we can design products to control these perceptions. The following is an overview of this related work.

Psychological aspects of driving motor vehicles, such as the perception of a vehicle's capabilities and its surroundings, have been shown to change driving behaviour (J.A. Groeger and Rothengatter 1998; John A. Groeger 2002) and driving safety (Chen, Haas, and Barnes 2007). People may change their driving behaviour based on the perceived risk of the surrounding environment (Michon 1985), vehicle type (Eyssartier, Meineri, and Gueguen 2017), and their mood (Precht, Keinath, and Krems 2017). Vehicle controls, such as haptic accelerator pedals (McIlroy, Stanton, and Godwin 2017), or transmission choice (Blommer et al. 2017) can also affect driving psychology. I extend this research in vehicle control to robot teleoperation, investigating how to prime different perceptions of the robot, and if the priming affects teleoperation performance.

Research has investigated using psychology-based designs in software to influence behaviour. For example, the attention and perception literature has been leveraged to increase the saliency of potential points of interest during teleoperation (Teng, Kuo, and Tara 2013; Daniel J Rea, Seo, et al. 2017), and the addition of haptic reminders have helped users notice changes in on-screen displays (J. J. Young, Tan, and Gray 2003). Engagement with software or the motivation to use it has been improved with game-based techniques (gamification, Li, Grossman, and Fitzmaurice 2012; Antin and Churchill 2011; Hamari, Koivisto, and Sarsa 2014) such as the inclusion of scores and audio-visual rewards in software tutorials (Li, Grossman, and Fitzmaurice 2012). Some research has even used priming to affect a user's experience in virtual environments (Nunez and Blake 2003) by making the users read materials related to the virtual environment they were about to enter. I follow this line of research by investigating priming teleoperator perceptions of a robot's capabilities, and observing how it may affect operator behaviour, and perception of the robot.

Priming has also occasionally been studied in human-computer interaction more broadly. For example, in virtual reality priming has been used to explore and change how people act and perceive themselves in a virtual space (Banakou, Groten, and Slater 2013), improve their experience in it (Nunez and Blake 2003), or change how they perceive themselves in that space. Other examples include the use of subliminal priming to aid learning (Chalfoun and Frasson 2011), visual hints to aid performance in visual search tasks (Harrison et al. 2013; J. J. Young, Tan, and Gray 2003), or analyses of how experimental design choices can prime participants and impact results (Bradley et al. 2015). In much of this work, a core theme is that technology is in a strong position to take advantage of priming due to the flexibility offered by software and virtual interfaces.

Social Human-Robot Interaction has explored the use of priming, or a variant called framing (Westlund et al. 2016) in social interactions between people and social robots. Multiple approaches employing subtle shifts in language while talking about robots have influenced how personal (Coeckelbergh 2011) or human-like (Stenzel et al. 2012) people view or treat the robot, or robots in general (Thellman and Ziemke 2017). People will even subconsciously imitate robot speech patterns when interacting with a robot (Brandstetter et al. 2017), an effect called lexical entrainment that shares similarities to priming. Priming can also make people believe an autonomous robot is actually teleoperated (Tanaka et al. 2016). Experimental challenges have been noted for HRI, in that priming effects can be very sensitive to types of stimuli used for priming or the surrounding context of the priming environment, and that detecting changes can be difficult (Westlund et al. 2016). I complement this body of work by investigating the impact of priming methods in teleoperation.

Others have attempted to change an operator's perception of a robot to shape their teleoperation behaviour; while not presented as priming per se, these works use stimuli to evoke feelings and change behaviour. For example, it has been shown that operators perceive certain behaviours as safer if they are the ones doing it (Basu et al. 2017), or that operators can be influenced to better avoid obstacles by haptic feedback without being aware they are being influenced (Hacinecipoglu, Konukseven, and Koku 2013). These works demonstrate that there are numerous approaches to how we can influence teleoperator psychology. I complement the research by investigating how to shape perceptions about a robot by limiting a robot's abilities, or by presenting a robot differently without additional features, sensor capability, etc.; I aim to shape perceptions and behaviour only through changing expectations about the robot.

4.4 Novel Teleoperator Priming Techniques

I investigated how priming operators about a robot's capabilities can impact how they drive a robot and their perceptions of the robot. My approach is to design priming stimuli that instill beliefs into the operator about how *safe*, or *unsafe*, the robot is and thus change how they operate the robot (Figure 4.4), building on prior research suggesting how perceptions of safety may impact driving behaviour (J.A. Groeger and Rothengatter 1998; John A. Groeger 2002).

To achieve this, my priming strategy was to convey properties of the robot's driving ability

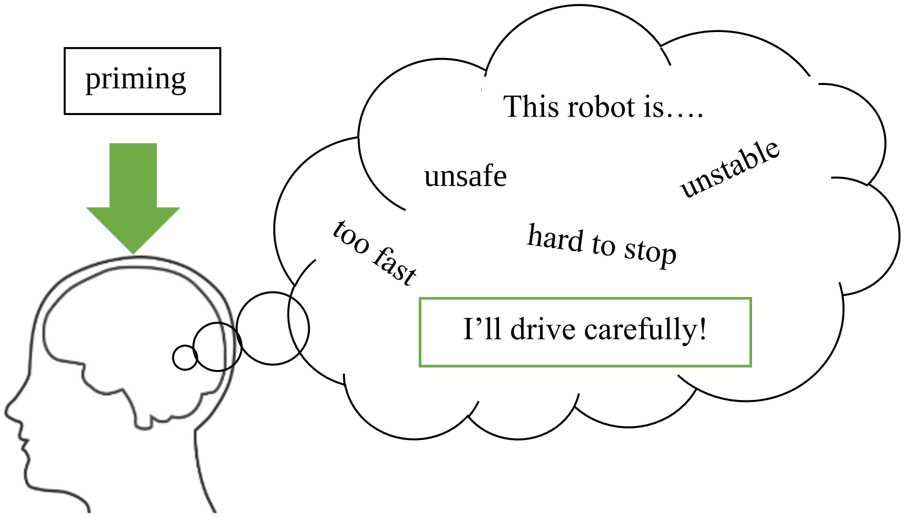

Figure 4.4. Our experiments test three different priming methods and observe their effects on an operator's driving behaviour and perception of their robot

relating to safety, such as how powerful the motor is, how easy it is to steer, and how durable the robot is to collisions. For each priming method, I developed three instances along a continuum, with one method suggesting *unsafe* robot characteristics, one suggesting *safe* characteristics, and one somewhere in the middle. In all cases, no actual properties of the robot or its response to commands (speed, ability, etc.) changed – in each case the same command (joystick pitch and yaw) created the same response in the robot. This deception enabled me to test the impact of the priming only.

It is not clear which approach – safer versus less-safe robot – would result in better driving. One could imagine operators would drive better when primed that the robot was *unsafe*, to compensate for the expected poor performance, and drive worse with the *safe* robot as they feel less pressure to be careful (Jonah, Thiessen, and Au-Yeung 2001). Inversely, perhaps the impression of *safe* or *unsafe* would encourage them to act likewise, where simply thinking about safety (or lack of) may make the person drive safer (or less so). As such, my hypotheses on the impact of priming on operator performance is non-directional: I do not hypothesize up-front what the impact will be.

I explored three different approaches: driving profile (change in robot performance implying robot physical ability), tangible priming (a continuous, physical indicator of robot ability), and descriptive priming (a verbally and visually explained, cognitive indicator of robot ability).

4.4.1 Driving Profile Priming

This method investigates the effects of changing a robot's top speed and acceleration (we term this a *driving profile*) on operator perception of the robot and task performance. Driving psychology has found that people operating a motor vehicle perceive the vehicle differently due to various factors, such as speed (Recarte and Nunes 1996) or sound (Kubo et al. 2004). Intuitively, when one drives a heavy or large car, the inertia causes acceleration and braking to feel more resistance than with a small or light vehicle. This gives the driver physical cues as to how a vehicle responds to commands, helping a driver make better choices and potentially leading to safer driving. This approach is used in racing video games to give different physical impressions of their virtual cars, as outlined in Section 4.1. I propose that similar effects may occur when teleoperating robots – e.g. a robot that slowly accelerates and brakes may *feel* heavier and thus, perhaps, more dangerous to operate. Simple software changes can simulate such physical differences in a robot; I investigate how artificially changing a robot's driving feel, to be heavier or lighter, could impact teleoperation.

I draw from this approach in video games, overlaying an artificial software control profile overtop of a robot's physical capabilities. I specifically investigate how changing a robot's maximum acceleration and speed in software – the *feel* of a robot – impacts operator performance, operation safety, and operator perceptions. Drawing from knowledge in psychology that found speed can affect perception of a vehicle (Recarte and Nunes 1996), I designed three driving profiles: default robot, acceleration-limited robot, and speed-limited robot.

For the acceleration-limited robot, I aim to create a sense of more mass in the robot by applying a simple limiter on the robot acceleration and deceleration. I anticipate that this may elicit safer driving behavior, given that the operator knows that it is more challenging to correct the robot's movement and to stop. In my implementation, I limit the acceleration of the robot by half of its default, and consider it the "unsafe" condition.

For the speed-limited robot, I simply limit the robot's speed, without modifying the acceleration profile, inspired by a smaller vehicle that may have a lighter engine and so cannot move as quickly. I anticipate that this robot would feel lighter and less capable, and as such, perhaps would cause less stress as an operator may feel there is less risk to the robot's surroundings. Further, I anticipate that the lower speed would negatively impact task completion time. My implementation limited the top speed of the robot to half of the default maximum speed, and I consider the speed-limited profile the "safe" condition.

The default robot keeps all settings unchanged from the manufacturer's provided settings. I assume that a default commercial-product setting would be a tested and reasonable control profile for the robot and task. Additionally, the default serves as a comparison point for the other two driving profiles, and is considered the "medium safety" condition. The robot's maximum speed is 2.6 km/h, which it accelerates to in approximately two seconds.

4.4.2 Tangible Priming

My strategy for tangible priming is to convey robot ability through the tangible response of the control method. My hypothesis is that a control method that takes more effort to use would convey a sense of a heavier, slower robot, which is safer to drive. Conversely, a control method that requires little effort to use would convey a lighter, faster robot, that may be easier to crash and break. This approach is inspired by force-feedback video game hardware such as joysticks or racing wheels that can provide physical and tangible feedback to the player during use.

This tangible method fits the model of priming where a constant, ongoing stimulus is provided (see Related Work, examples such as Aarts and Dijksterhuis 2003; Ackerman, Nocera, and Bargh 2010). Instead of a single, up front priming stimulus, my tangible priming method continuously reminds the operator of their experience and prior knowledge which in turn may continuously evoke a priming effect (Bargh, Chen, and Burrows 1996).

Specifically, I use different spring stiffnesses of the joystick used to drive the robot to impart this tangible feel. I use three static settings for joystick stiffness, one per condition: high stiffness (to convey *safe* robot), low stiffness (to convey *unsafe* robot), and a mid-point in between. Note that the stiffness is fixed per condition (static) and did not change during operation. In all cases, the joystick stiffness provides a constant reminder of the robot's ability.

A key element of priming is how the technique is introduced to operators. I simply tell people that "each robot will interact with the joystick differently, based on the robot's physical

Figure 4.5. The joystick we used for tangible priming – Microsoft Sidewinder Force Feedback 2 USB joystick.

design." my goal is to avoid telling people what my intended impact was (e.g., safe vs. unsafe robot), and to let them know the joystick changes are intentional and relate to robot capability.

We implement this technique using a force-feedback joystick (Figure 3), which has a programmable stiffness setting. I use 100% of device maximum spring strength and friction for the safe condition, 10% for the unsafe condition (0% would not provide enough stiffness to return the joystick to a neutral center position), and 50% for the middle case. The strongest setting (*safe*) takes noticeably more force to operate than a regular joystick but is not onerous to operate and I do not anticipate fatigue to be an issue. The weakest setting (*unsafe*) is strong enough to automatically return to a centered position after being pushed but puts very little force onto the user. The robot response to a given joystick input (pitch and yaw values) does not change: a given joystick position will result in identical behaviour regardless of stiffness settings.

4.4.3 Descriptive Priming

For this method I investigate if priming by altering how I describe a robot to an operator will impact perceptions of robot ability after operating the robot. I simultaneously employ verbal description and visual aids (Figure 4.6) that explicitly define specific robot performance characteristics, while secretly the robot performance does not change.

I achieve the impression of safety by describing four robot characteristics, selected as attributes that I expect non-expert operators to easily understand and relate to operation safety. These are robot "balance," "toughness," "motor power," and "traction." I present this information on paper (Figure 4.6), along with a scripted explanation for introducing each robot and variable that emphasizes the safety and risks of each, but without explicitly telling operators my purpose. The labels and the descriptions I use are described in Table 4.2.

I tell operators that these measures are derived from a number of components in the robot, as rated by the manufacturer, and I further give the robot names to suggest their safety level (Figure 4.6). People keep the relevant specification sheet in front of them during operation.

4.5 Investigating Priming: Experiments

I conducted three studies to investigate the impact of each of my priming methods on teleoperation (reported in Rea, Hanzaki, et al. 2017;Rea and Young 2018). Specifically, I investigated the impact of priming on operator perception of a robot and teleoperation performance, after a time operating it. While each study followed a similar procedure, I do not analyze this as a single study (with priming method as a between-subjects variable) given that each study presented results that inspired the next priming design and exploration. I further introduced minor changes between each study, explained below. Thus, I present and analyze my results more accurately as a series of separate studies.

We conducted within-participants studies where a single participant completed a task with all three robot conditions (safe, unsafe, in-between). A within-participants design enabled participants to directly compare and contrast the robots between priming conditions, and

Table 4.2 A list of the descriptively primed robot properties and how we explained them to participants

Name	Description: "A robot's ability to …"
Balance	stay upright easily, regardless of surface, obstacles, or operation
Toughness	not be damaged from collisions
Motor Power	accelerate quickly to a high top speed
Traction	turn quickly and safety
Battery	continue operating for long periods of time

Figure 4.6. The priming sheets were explained and given to operators before using the robot. The sheet of the robot being driven was left on the desk for the operator's reference. a) the safe condition b) the middle condition c) the unsafe condition

further provided more statistical power by factoring out individual differences in driving ability, susceptibility to priming, etc.

A key element of my study design was to give people a representative experience operating the robots; particularly for the descriptive cases, participants need enough experience so that they do not simply report back on what they were told. Ostensibly, after driving each robot for a period of time I could reasonably expect participants to notice the differences between the robots (or lack thereof), despite the priming stimulus.

4.5.1 Task

We tasked participants with navigating a telepresence robot through an obstacle course (Figure 4.7). They were instructed to drive and complete the task as quickly as they felt comfortable, while trying to avoid colliding with obstacles, walls, etc. As each participant completed three conditions (safe, unsafe, in-between), I marked three paths through the course with arrows to indicate where to turn (path order was counterbalanced across conditions), with all paths having a similar difficulty: same number of turns and distance (Figure 4.8). Each path took approximately 2~5 minutes per lap, depending on driving speed, the number of collisions, and

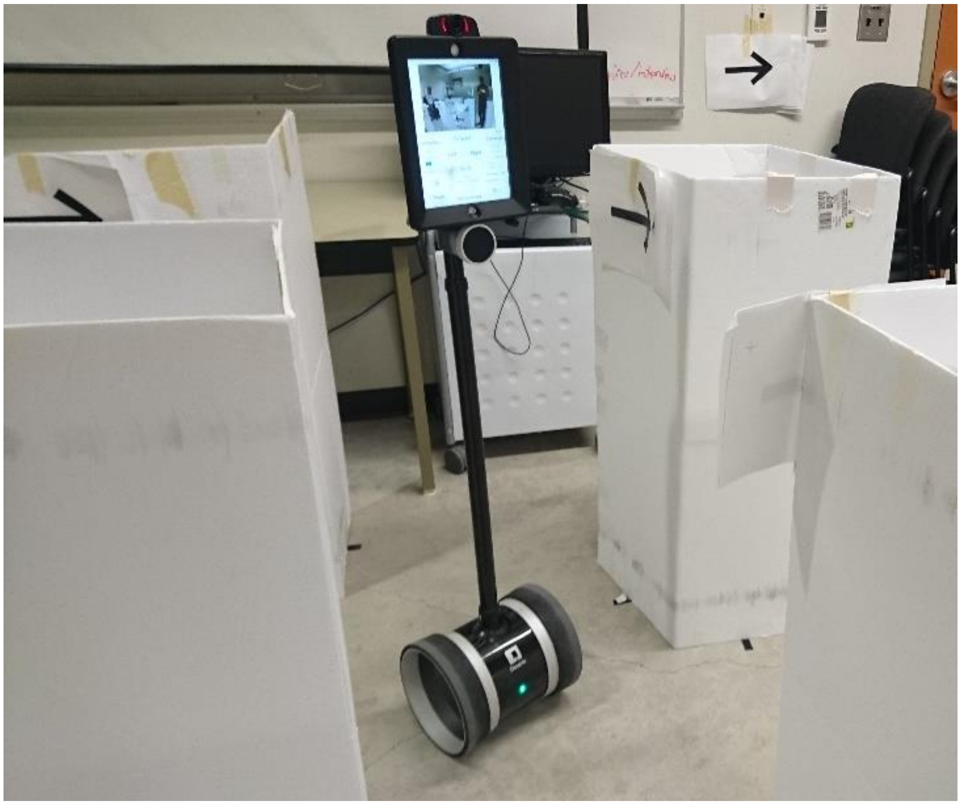

Figure 4.7. A robot is driven through an obstacle course. We primed operators to believe that they were driving robots with different capabilities and potential risks. We examined how priming changes teleoperation behaviour and perception of the robots.

overall participant skill. For each condition, participants were asked to first complete a training lap, followed by two laps for the study.

4.5.2 Instruments

Participants operated a Double 2 robot (Double Robotics) with a 150 degree field-of-view camera. The robot's camera feed (640 x 480 pixels) was viewed full-screen (with black bars on the wide-screen sides) on 24-inch monitor, in a separate space from the robot, and participants were seated at roughly the same position with respect to the monitor. The system maintained at least 15 frames per second, but as fast as 30 frames per second, depending on network health.

While driving the robot, participants wore headphones that relayed sound from a microphone mounted on the robot in the remote space. Participants used a Microsoft Sidewinder USB Force Feedback 2 joystick with stiffness set to 10%, 50%, and 100% for the tangible study, and 50% for the descriptive and driving profile studies.

Participants completed questionnaires (detailed in the next subsection) on a separate monitor using Google Forms.

4.5.3 Measurements

My performance measurements were selected as simple teleoperation measures used in prior work (Chen, Haas, and Barnes 2007; Daniel J Rea, Hanzaki, et al. 2017) – completion time, collisions, and perceived workload. I additionally measured teleoperator perception of the robot and its physical capabilities.

Pre-experiment, I gathered demographics information to better understand the variance in my sample. I collected information including age, gender, frequency of playing video games, frequency of driving, and self-reported driving skill.

For each condition, a researcher in the room with the robot measured completion time and collisions. Perceived workload was measured post-condition with the NASA TLX (Hart and Staveland 1988) self-report questionnaire. To get a sense of a participant's perceptions of a robot's capabilities (and the effects of my priming) I also administered 5-point Likert-like scale items inquiring about a participant's opinions on the robot's speed, weight, steering, durability, power, safety, and responsiveness. Participants then completed free-form written questions inquiring about their experience. These questions were optional, and asked participants for any positive, negative, or other feedback they wished to provide me about the robot and teleoperation experience.

4.5.4 Procedure

A similar procedure was followed for all studies, with differences highlighted in the corresponding sections below. Participants were first given a briefing of the experiment and signed an informed consent form. Participants were told that they will test 3 new prototype telepresence robots in order to help me evaluate the safety and drivability of each robot for new users, but were not told specifically about the robots being designed for different safety levels. I described the robots as being similar in size and shape, but with different internal components that may change how they perform. I explained the overall procedure of the experiment, and introduced the joystick and obstacle course. Further, before starting, I explained either the connection between the robot and joystick (for tangible priming), or a high-level overview of the robot data sheets (for the descriptive priming), as explained in my priming method overview (Section 4.4). The participants were seated in a room separate from the robot and obstacle course.

Following the introduction, each participant completed the task three times, once per priming condition (safe, middle, unsafe), with the order of the priming conditions and the path through the course (Figure 4.8) fully counterbalanced. Before starting each of the three conditions, participants were first asked to complete a training lap, before the two laps of their task. This training allowed participants to become familiar with the new obstacle course (and reduce confusion from the new course) and gave them additional practice with the robot. In addition, in the descriptive and tangible priming methods, this practice added to my priming story – participants believed they were operating a new robot, and I told them that the training was for them to get used to the differences between each robot model.

After each of the three conditions, I administered the post-condition questionnaires described earlier (NASA TLX, perception of robot abilities). To transition between conditions, I disconnected the robot from the control interface to give the illusion of switching to a new robot, although the same robot was reconnected upon starting the new condition (in the driving profile study, the software that controlled the robot's top speed and acceleration where changed at this point as well). As the participant was in a space separate from the robot, I was able to maintain this illusion.

Post-test, participants were debriefed about the priming purpose and in the tangible and descriptive priming case I explained the deception (that it was a single robot only, with no real differences). The experiment was then re-explained in the context of the deception and how the deception helps achieve the research goal. The participants were encouraged to engage with a discussion with the researcher about the experiment. My university's research ethics board approved all studies.

4.5.5 Study: The Effects of Driving Profile on Teleoperation

Our experiment for studying the effects of change a robot's driving profile had small methodological variations from the other studies in this chapter. Specifically, I aimed to measure operator experience, and after each condition I added 3 extra TLX-like scales after the NASA TLX that inquires into operation Enjoyability, Confidence, and Perceived Weight

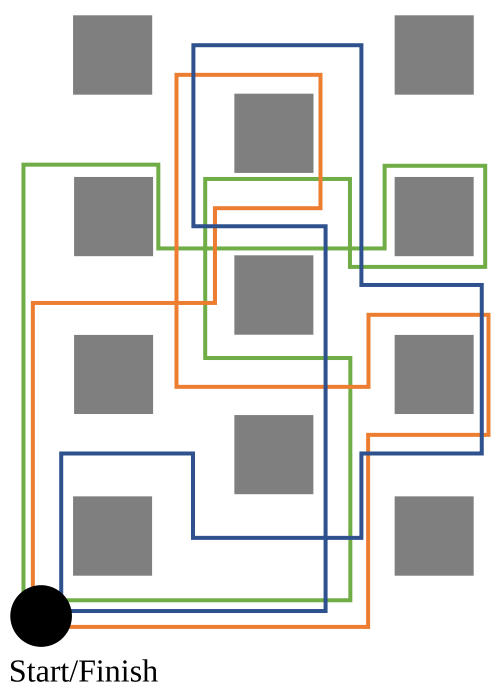

Figure 4.8. The room and obstacle layout used in the study design, with the three paths through it.

(light to heavy). These were specific to this variant and so I did not include them in the other conditions. As the robot's actual abilities were changing, I did not distribute the scales recording operator's perception of robot speed, weight, steering, durability, power, safety, and responsiveness as I do in the other experiments.

Participants performed a similar obstacle-avoidance task. However, for this experiment only, I had to create a new layout for paths and obstacles. I kept the same design goals: each path should have a similar length and number of turns and total distance travelled (Figure 4.9). Obstacles were the same size, but placed slightly farther apart (60 cm), and each path was marked with coloured arrows.

Participants were briefed that I was testing three telepresence robots that were identical in all ways except their motors, and that they would be helping me select the one that was easiest to control – this deception (that I changed the motors, not the software) was to encourage the idea the robots may feel differently the way vehicles with different engines may feel different. I otherwise followed the procedure outlined in Section 4.5.4.

The order of which path the conditions used was fixed (red, then green, then blue, Figure 4.9), but the order of the driving profiles was counterbalanced between participants.

Results

I conducted the experiment with 19 participants (8 female, average age of 27) recruited from my local university campus. I performed repeated measures ANOVAs on TLX sum, number of collisions, and time to completion.

Outlier analysis indicated three participants being at least 1.5 times the inter-quartile range of my data in at least two statistical tests. In addition, I observed that these participants demonstrated poor spatial awareness while driving the robot, for example, repeatedly getting the robot stuck without noticing, resulting in repeated similar collisions and long completion times, uncharacteristic of other participants. I excluded them from my analysis as outliers, resulting in n=16.

I found an effect of driving profile on completion time, ($F_{2,30}$=8.2, partial η^2 =.35, p<.001).

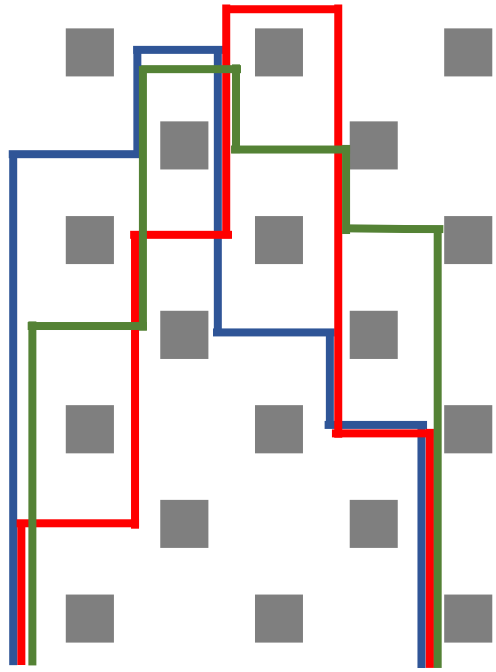

Figure 4.9. The three courses (labelled red, blue, or green). These paths were used in the driving profile experiment only.

Marginal means are shown in Figure 4.10c. Bonferroni corrected post-hoc tests showed the default profile was faster than the acceleration-limited profile (mean difference = 42.8 seconds, p<.01, 95% CI [15.0 seconds, 70.5 seconds]).

We found an effect of driving profile on number of collisions, ($F_{2,30}$=6.2, partial η^2=.29, p<.01). Marginal means are shown in Figure 4.10b. Bonferroni corrected post-hoc tests showed the speed-limited profile had fewer collisions than the acceleration-limited profile (mean difference = 6.4 collisions, p<.01, 95% CI [1.5 collisions, 11.2 collisions]).

I found an effect of driving profile on subjective workload, ($F_{2,30}$=4.4, partial η^2=.23, p=.02). Marginal means are shown in Figure 4.10a. Bonferroni corrected post-hoc tests showed the speed-limited profile was less demanding than the acceleration-limited profile (mean difference = 16.4 points, p=.01, 95% CI [3.1 points, 29.8 points]).

To better understand the effects of each condition on different types of workload, I performed a post-hoc ANOVA across the 6 individual TLX scales with Holm-Bonferroni correction, a standard practice with the TLX (Hart, Sandra 2006). I found a trend of driving profile on temporal load, ($F_{2,36}$=4.6, partial η^2=.20, p=.068). Default had the highest temporal load (mean=12.0 points, 95% CI [8.7 points, 13.0 points]). Acceleration-limited had second highest temporal load (mean=12.4 points, 95% CI [9.1 points, 13.6 points]). Speed-limited had the lowest temporal load (mean=8.9 points, 95% CI [6.8 points, 11.1 points]).

I found an effect of driving profile on perceived performance, ($F_{2,36}$=4.6, partial η^2=.25, p<.04). Note higher scores mean worse perceived performance. Acceleration-limited had the worst perceived performance (mean=10.8 points, 95% CI [8.8 points, 12.9 points]). Default had second worst perceived performance (mean=10.6 points, 95% CI [8.6 points, 12.6 points]). Speed-limited had the best perceived performance (mean=7.2 points, 95% CI [5.2 points, 9.1 points]).

I found an effect of driving profile on perceived effort, ($F_{2,36}$=5.6, partial η^2=.24, p=.04). Acceleration-limited had highest frustration (mean=10.3 points, 95% CI [7.7 points, 12.8 points]). Default had second highest frustration (mean=8.5 points, 95% CI [6.1 points, 10.8 points]). Speed-limited had the lowest frustration (mean=7.3 points, 95% CI [5.1 points, 9.4 points]). All other tests were not significant.

Perception and Robot Capability

Overall, my speed-limited robot profile had the strongest performance; it had the lowest subjective workload, least number of collisions, and second fastest completion time. Surprisingly, the speed-limited robot had only a 10% slower completion time than the quickest (the default profile), even though it was limited to 50% of the maximum speed of the other two profiles. Interestingly, I could not detect a difference between the speed-limited and default profile's completion time. Though I did not perform equivalency testing, this result suggests that it is unlikely a large effect exists between the two profiles.

The acceleration-limited robot performed poorly. Participants perceived it as having the highest workload, they hit the most obstacles with it, and completed the courses in the slowest times. In each measure, post-hocs found at least one profile statistically better performing than acceleration-limited.

The default profile was only the best performer in terms of completion time, although it was not found to be statistically different from the speed-limited profile. Interestingly, in both collisions and subjective workload, I could not statistically distinguish it from either of the other profiles. While it is likely still the middle performer in these measures, default's measured means were closer to the acceleration-limited profile's performance, rather than the speed-limited profile.

When I looked deeper into the individual TLX scales, I found that the speed-limited profile was perceived as requiring less effort to pilot (agreeing with the overall workload result), and achieved the highest perceived performance. Better perceived performance is interesting, as the speed-limited profile was 50% slower and had a negative impact on completion times (10% slower) on average. I found a trend in perceived temporal load which, if confirmed with more data, would imply that people may have felt less rushed, even with the slower speed. This suggests that people may perceive collisions as a more stressful occurrence than slow movement, even when my obstacles where harmless cardboard. It is possible this effect could be stronger if real people are around the robot.

While my results imply that the acceleration-limited robot performed badly, I note that there should be a relationship between all three of my measures: collisions take time to recover from, so perhaps the high number of collisions increased the completion time, and the stresses from both these factors contributed to a worse perceived workload. Less crowded and collision-prone environments than ours may result in the acceleration-limited profile performing differently.

My results suggest that, in some conditions such as my crowded setup, a speed-limited robot can help operators avoid collisions without a large increase in lost time. Further, I found evidence that operators may perceive collisions with obstacles to be more stressful than a slow robot speed. How this result generalizes to different tasks and robot speed configurations remains important future work. These results suggest that it may prove useful to investigate dynamic speed limits placed by software, based on the surrounding environment.

Reflecting on the impact of driving profile as a priming method, I realized I could not attribute the performance differences to the driving profile influencing the perception of the robot. My observed differences could have instead simply been a result of the robot's actual different

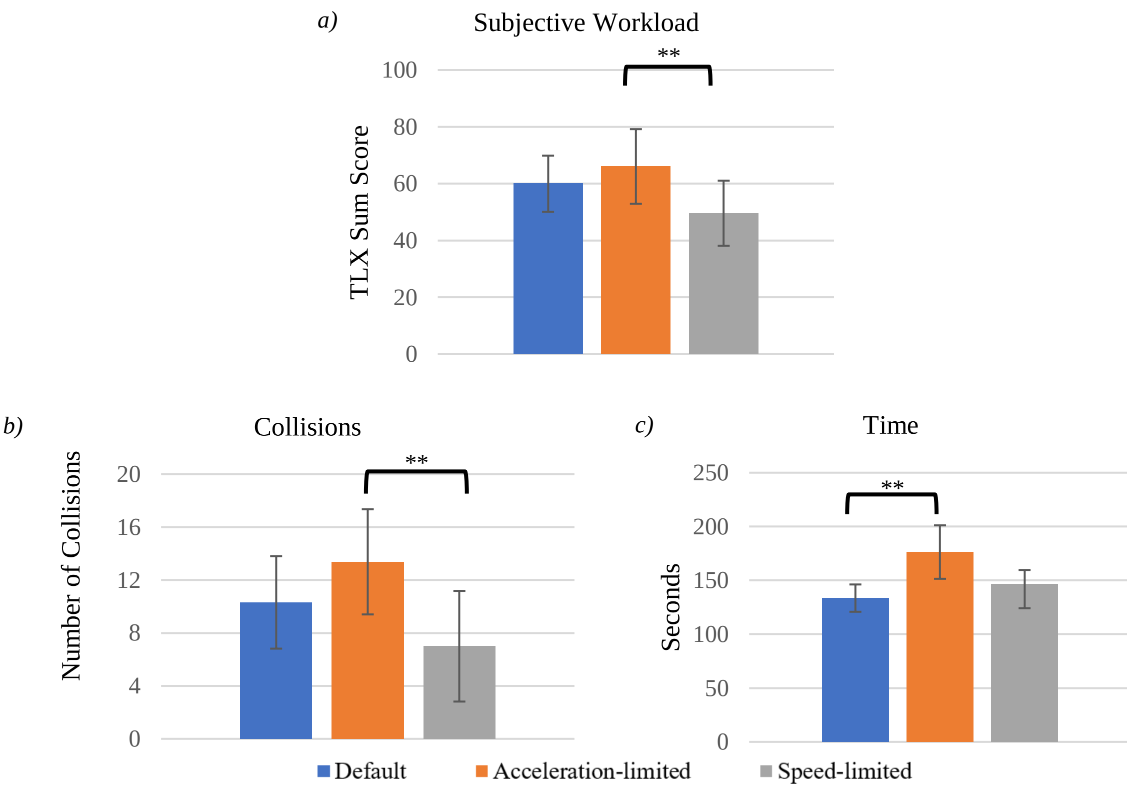

*Figure 4.10. ANOVA results; error bars show 95% confidence interval. We found main effects of profile on all measures a) Workload (range [0,120]): speed-limited was less demanding than acceleration-limited. b) Collisions: speed-limited had less collisions than acceleration-limited. c) Time: default was faster than acceleration-limited (** is p<=.01).*

performance (being faster, easier to stop, etc.). Although my results support my approach of learning from video game techniques to influence operator behaviour, I cannot make strong conclusions regarding the priming component.

4.5.6 Study: Tangible Priming

For the tangible priming study, I recruited 25 participants; however, one did not complete the experiment due to technical issues. Two other participants were identified as outliers: I observed them not attempting to avoid obstacles (e.g., laughing and pushing obstacles around seemingly on purpose), and this was reinforced from their data (>1.5 inter-quartile range). This resulted in 22 participants (mean age of 24, standard deviation of 6.3 years; 12 female).

Results

To investigate whether the tangible priming worked, I conducted Friedman's ANOVA tests on my Likert-like scale perception data. I found statistically significant results for perceived speed, perceived steering ability, perceived durability, and perceived safety (Table 4.3). Other tests on perceived teleoperation experience were not significant. Further, I found no effect of variables from the demographics questionnaire (video game, driving experience) on any of my measures.

Both completion time and number of collisions were right skewed (non-normal, Shapiro-Wilk test, $p<.05$), and were corrected using a square root transform.

To investigate teleoperation performance, I performed repeated-measures ANOVAs on completion time, collisions, and perceived workload. I found a statistically significant, medium effect of tangible priming condition on number of collisions ($F_{2,42}=5.2$, $p=.01$, $\eta^2=.20$, Figure 4.12). Post-hoc tests (Bonferroni familywise correction) found the safe condition to have on average 4.8 fewer collisions (42% fewer) than the unsafe condition ($p=.001$, 95% confidence interval of the mean difference [1.8 collisions, 7.8 collisions]).

We further found a statistically significant medium effect of tangible priming on perceived workload (NASA TLX sum, $F_{2,42}=3.6$, $p<.04$, $\eta^2=.14$, Figure 4.11). Post-hoc tests (Bonferroni familywise correction) found the non-safe condition to have on average 5.0 points higher (14% higher) perceived workload than the safe condition ($p<.04$, 95% confidence interval of the mean difference [.22 TLX points, 9.7 points]).

Table 4.3. Mean ranks and chi-square values for perceptual effects for tangible priming. Higher ranks for steering, durability, and safety are considered "better" and higher ranks for speed are considered "faster". Note that safe is considered the slowest with better steering and durability than unsafe, with middle being a mix of the two. All listed values are p<.05. Omitted variables are n.s.

	unsafe	middle	safe	$\chi^2(2)$	p
speed	2	2.4	1.7	7.0	.03
steering	1.6	2.2	2.2	6.6	.04
durability	1.7	2	2.3	6.9	.03
safety	1.7	2	2.3	8.0	.02

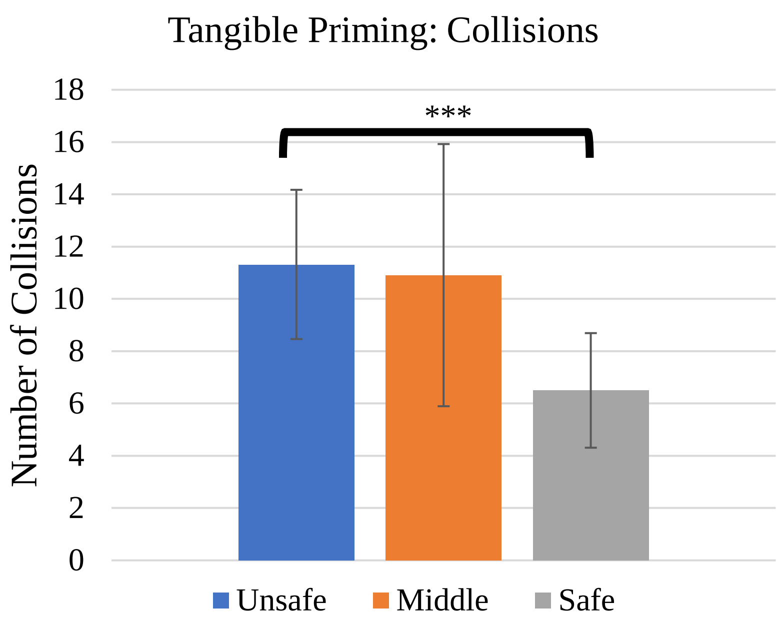

*Figure 4.12. Average collisions per condition. ***p<.001. Error bars show 95% confidence interval.*

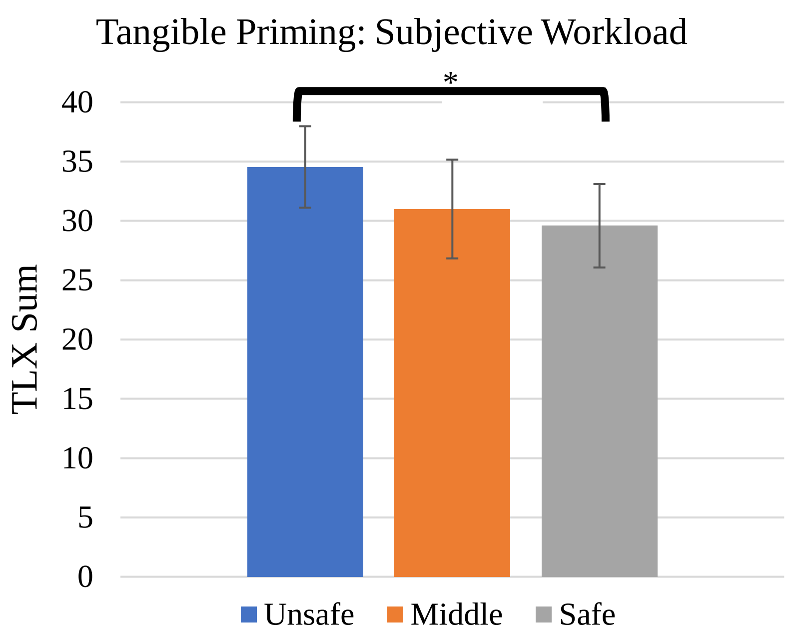

*Figure 4.11. Average TLX sum score per condition. *p<.05. Error bars show 95% confidence interval.*

Discussion of Tangible Priming

Our results indicate that my tangible priming conditions caused participants to perceive the robot and teleoperation experience differently: I found differences in perceived safety, durability, steering ability, and speed. Further, these differences aligned with the expected impact of my specific priming strategy. Given that the robot reacted and responded identically in all conditions, and participants spent time controlling the robot, if the priming was not effective it would be reasonable to expect participants to rate the robots based on how it actually performed, and perhaps notice that the robots were the same (or only slightly different). However, the fact that participants rated the robots differently despite this is a clear indication that the tangible priming method worked to shape participant perception of the robot and teleoperation experience.

I further found a significant difference in collisions, with the non-safe condition having a 42% reduction (average 11.4 in the unsafe, and 6.6 in the safe), and participants reporting lower task load with the safe condition (average 5.0 TLX points, 14%, lower than the unsafe condition). I cannot speak to the exact mechanism by which my tangible priming method may have caused this improvement in driving: perhaps the priming encouraged people to drive more slowly, take fewer risks, or take wider turns around obstacles. Further study is needed to understand the specific mechanisms and how they produce the effect.

Looking at my performance and perception results together, I see that people drove the safe condition in a safer manner and perceived it as safer than the other conditions. While some related work suggests people may drive a safer vehicle more recklessly (Jonah, Thiessen, and Au-Yeung 2001) I reemphasize that, in my specific implementation, I had plausible explanations for either an increase or decrease in safety and thus did not hypothesize a specific direction of effect (see my priming technique overview). Regardless, my priming method was a success, considering the changes in perception (e.g., of speed or steering capabilities) when participants drove an identical robot each time. I conclude that the physical properties of an input method can be used to prime users and change their perceptions of the robot and may also impact their performance.

I note, however, a potential confound in the study: the *usability* of the different stiffness settings may explain the performance difference. That is, perhaps the stiffer joystick was simply easier to use than the looser setting, explaining the reduced collisions, and thus the improved perception of safety. I re-visit this issue, and present the results from a follow-up study, later in this chapter.

4.5.7 Study: Descriptive Priming

I recruited 24 participants (none participated in the Descriptive Tangible Priming study); three were removed as outliers as they did not attempt to avoid obstacles (e.g. driving full speed and not stopping for any obstacle) or did not appear to understand the instructions (e.g., frequently took wrong turns in the obstacle course). This was reinforced as outliers in the data (>1.5 inter-quartile range). This resulted in 21 participants (mean age 24, SD 6.3 years; 12 female).

The priming specification sheets (Figure 4.6, page 127) were explained in detail to participants at the introduction of the study, and the sheet associated with each condition was left with the participant during the task. Participants were given time to review the specification sheet (the priming) before each condition, and the sheets were removed during the post-condition questionnaire.

In the tangible priming study, I noticed a subjective improvement to participants' performances as the study went on, due to, I presume, becoming more skilled at operating the robot. While this improvement was mitigated somewhat in my results due to counterbalancing and initial training lap, to further reduce potential learning effects I added an additional up-front training step after the initial explanation, and before the first condition: participants practiced using an additional, similar path through the obstacle course for two laps. Participants were told they were piloting the current commercially available robot model (compared to the "prototypes" that followed).

Additional self-report measurements were added post-experiment to reflect the details of my priming. Participants rated the robots on the criteria I used in the priming specification sheets (Figure 3), asking what their impression was of the robot's motor power, traction, balance, toughness, and battery life. Participants were specifically asked to report based on their teleoperation experience, not on their memory of the specification sheets. This final questionnaire was completed on paper.

Results

To investigate whether the priming worked, I conducted Friedman's ANOVA tests on my post-condition Likert-like scale data. I found statistically significant results for perceived speed, perceived steering ability, perceived durability, and perceived safety (see Table 4.4). Other tests on perceived teleoperation experience were non-significant. Friedman's ANOVA tests on the post-experiment specification sheets found statistically significant results for balance and motor power, with trends for toughness and traction. These results are also included in Table 4.4.

With repeated measures ANOVAs I found no significant results on completion time ($F_{2,38}$=.2, p=.83, η^2=.01, means for unsafe=165s, middle=176s, safe=171s), collisions ($F_{2,38}$=.2, p=.68, η^2=.01means for unsafe=6.0 collisions, middle=5.4 collisions, safe=5.8 collisions), and perceived workload ($F_{2,38}$=.7, p=.48, η^2=.04, means for unsafe=29.4 points, middle=29.4 points, safe=27.7 points).

Qualitative results

Given the lack of impact of description priming on teleoperator performance, I performed post-hoc open-coding qualitative analysis on participant short-form responses to learn more about operator driving experience. Coding was done with a single coder with thematic analysis; the purpose of this analysis was not to make definitive conclusions about why participants acted in a given way, but to better understand how and why participant's may have rated the robot's perceived abilities differently, to inform follow-up work.

I found that 20 participants (83%) made explicit comparisons between the robots' capabilities and their teleoperation experiences with them:

> I love the response time and the power of the [unsafe condition]. It's quicker than the [safe condition] and I felt like the wind. – p9
>
> I felt more in control with [the safe condition] – p19
>
> Aside from durability, everything else about [the middle condition] felt more stable – p14

These comments covered a range of aspects of teleoperation, which I found to reflect consistent opinions of a robot's perceived abilities across conditions. Further, these comments aligned well with the primed robot characteristics.

All eight participants who mentioned speed wrote that the unsafe condition was faster than other robots:

> It's quicker [unsafe condition] than the previous robot and I felt like the wind – p33
>
> It was hard to keep the balance on this robot [unsafe condition] as it was light and had more speed. – p16

Speed was less commonly mentioned in the other conditions (three times total), which were characterized as slower:

> [The middle condition] didn't accelerate as fast as the other robots – p2

Control was another common theme, where six people reported the safe condition as having better control:

> I liked how in control I felt of the steering and acceleration. There were no surprises. – p11

There was one comment to the contrary. In contrast, three people mentioned that the middle condition had better control than the unsafe condition, and two mentioned that the unsafe condition had worse control overall.

Finally, "responsiveness" was another common theme. The unsafe robot was most commonly discussed, with seven participants saying that it was more responsive, for example:

Table 4.4. Mean ranks and chi-square values for perceptual effects for descriptive priming. Omitted tests are n.s.

	unsafe	middle	safe	$\chi^2(2)$	*p*
speed	2.5	1.5	1.9	8.6	.01
weight	1.6	2.1	2.3	6.5	.04
power	2.4	1.7	1.9	7.5	.02
safety	1.6	2.3	2.1	8.3	.012
balance	1.6	2.5	1.9	12.7	<.01
motor power	2.4	1.6	2.0	7.4	.03
toughness*	1.7	2.0	2.3	4.6	.10
traction*	1.6	2.1	2.2	4.9	.09

> It responds quickly, and seemed to navigate at relatively high speed. – p13.

The four participants who mentioned responsiveness with the middle safety robot all had comments similar to:

> The robot felt more flimsy and unresponsive – p11.

Only two participants mentioned the responsiveness of the safe condition. One participant mentioned it was “more responsive” – p22, while the other disagreed:

> The robot is slower, doesn't have a faster response rate, motor power is definitely weak. My head is hurting trying to operate this robot. – p9

Discussion of Descriptive Priming

In this experiment, I investigated the impact of priming teleoperation operators using a visual and verbal description of the robot. My results suggest that descriptive priming (using paper and speech only) successfully changed participant perception of the robot, and their experience teleoperating it, even after operating it for some time. I successfully altered participant perception of robot speed, weight, power, and overall safety. Further, my post-test questionnaire results indicated that my non-safe condition was successfully primed to be seen as riskier than my safe condition in terms of balance and motor power, with trends pointing to potential priming in toughness and traction. These results emerged despite participants driving the exact same robot in each condition.

My qualitative results further supported this and highlighted the effectiveness of my priming. More than simply memorizing the details provided to them, the conviction and tone in the written feedback suggests that the participants may have believed that the differences were real, despite having operated the exact same robot through a task repeatedly.

I did not find any performance change in terms of completion time, collisions, or perceived workload. It is possible that there is still a small effect that went undiscovered due to my small sample size of 21. If there is indeed no effect on performance, it will be important to further investigate how this disparity between perceptions and performance can happen, and what it means for long-term use. Importantly, my results suggest that I can improve user perception of the safety or physical capabilities of the robot without sacrificing performance or changing functional aspects of the design.

4.5.8 Reflection on Tangible and Descriptive Priming

Both priming methods were effective at changing the user's perception of the robot, while the actual experience of driving the (identical) robots did not seem to counteract the priming. That is, even after driving the (identical) robots themselves for multiple trials and training, for upwards of 30 minutes, participants rated the robot capabilities differently, but similarly to how I primed them. Both methods primed changes in perception of a robot's speed and safety, but there were differences in perception of the robot between the two methods: tangible priming changed perceived steering and durability, and descriptive priming changed weight and power. While this makes sense for the descriptive priming case – it matches my priming focus – for the tangible case the connection to steering and durability is less clear. Further, I observed a difference in actual driving performance for tangible priming, with the stiffer joystick (safe priming) resulting in, on average, 4.8 fewer collisions than the looser joystick (unsafe priming). This highlights the need to consider and the technique used to prime, and how choices may inherently work well for some perception and behaviour outcomes and not others. For example, considering the quality and weight of the paper used in descriptive priming, analogous to prior results showing the importance of the quality of paper and a clipboard (Ackerman, Nocera, and Bargh 2010).

It is worth considering further why only the tangible case impacted driving performance. First, I note that the tangible condition also resulted in a difference in operator perceived workload, with the safe condition resulting in a 14% reduction (in TLX score) compared with the unsafe condition; no difference was found on workload with descriptive priming. Perhaps one reason is that the tangible priming is directly linked to control (being the joystick) while the description is more abstract. Or, perhaps this is due to the tangible priming being a more salient constant reminder of the priming in comparison to the descriptive paper which just sat beside the joystick, while the participant was busy with the task. These points require further study.

Another possibility is that the impact on driving performance may not have been due to the priming. Specifically, perhaps the joystick stiffness itself has a usability impact, where one joystick (in this case, the stiffer one) is simply easier to control than the other (the less stiff one). If that is the case, then it is the joystick usability – and not my priming method – which may be responsible for the driving performance and workload result. I conduct a follow-up study (detailed in the next section) to explore this possibility.

Overall, these two studies were a success. I was able to leverage priming to consistently change operator perceptions of the robot, perceptions which persisted even after using the robot for upwards of 30 minutes. While the impact on actual driving performance was mixed, I note that shaping perceptions itself is an important element of interface design (J. E. Young et al. 2009), as it can shape expectations, user workload or stress, and affect technology adoption on the long term.

4.6 Study: Joystick Stiffness – Priming or Usability?

I conducted a study specifically to test the usability component of my tangible priming method, which used joystick stiffness to represent robot capability. That is, I inquired about whether joystick stiffness impacts robot control sufficiently to explain my tangible priming results, where perhaps a stiffer joystick is simply easier to control than a looser joystick. Such a result would require me to re-analyze my results from my tangible priming study above, as it would perhaps be the usability of the joystick – not the tangible priming – the explains the improved driving performance I found with the tangible priming case.

My approach was to replicate my tangible priming study, while trying to remove the priming (and included deception), explaining clearly the joystick stiffness manipulation. Analyzing this alongside the results from the tangible priming study enables me to separate the effects of the joystick usability from priming effects. On the one hand, if I still find the same effects without the priming, then I can conclude that it was the usability – and not the priming – that explains my results. On the other hand, if I do not find an effect of the joystick stiffness on teleoperation performance, then this lends support to my conclusion that priming is the driver of my earlier results.

4.6.1 Procedure

I the same procedure as explained as the tangible priming experiment. The primary difference was that I did not tell participants that the joystick stiffness represented the robot capabilities (priming). I explicitly told them all the conditions of the study. That is, I told them that they are driving the same robot repeatedly, and that the only thing I change is the joystick stiffness. I further explained that, although the joystick stiffness changes, the response of the joystick does not change: a given joystick movement or position will result in the exact same robot reaction, regardless of stiffness setting.

Participants first completed the same pre-test demographics questionnaire, before being introduced to the system. All conditions were explained (as above), and participants completed three conditions, with the same three joystick stiffness settings used in the tangible priming study (with the same counter balancing). To maintain consistency with the original tangible priming study, the extra training session before the experiment (added in the descriptive priming experiment) was not included.

Each condition consisted of a training lap, followed by two laps that were recorded. During the condition I recorded completion time and collisions, and after each condition I administered the perception questionnaires from my tangible priming study. I re-emphasized that I was only changing joystick stiffness before each condition. Post-experiment, I elicited general qualitative feedback (as in previous studies), debriefed participants, and gave them an opportunity to ask questions.

4.6.2 Results

I recruited 18 participants (mean age 24, SD 9.4 years; 12 female) – none participated in the prior descriptive tangible or descriptive priming studies.

To investigate teleoperation performance, I performed repeated measures ANOVAs on completion time, collisions, and perceived workload; I found no significant results for any of the three variables (summarized in Table 4.5).

Table 4.5. ANOVA results for our three main performance measures with no priming.

	F-ratio	p	η^2	Not safe	middle	Safe
Completion time (s)	$F_{2,34}$=0.6	.56	.03	208	213	219
Number of collisions	$F_{1.4,23.8}$=2.9	.07	.14	6.3	7.2	8.2
Workload (TLX score)	$F_{2,34}$=2.2	.12	.11	32.1	28.3	30.0

To investigate if there were any priming effects on operator perception of the robot, I conducted Friedman's ANOVA tests on my post-condition Likert-like scale data. I found statistically significant results for perceived weight, perceived steering ability, perceived durability, and perceived safety (see Table 4.7). Other tests on perceived teleoperation experience, including perceived workload (NASA TLX), were non-significant.

4.6.3 Discussion – Priming or Usability?

I found no statistically significant impact of joystick stiffness on any measure of driving in this no-priming study, which contrasts the findings in my prior tangible priming study.

First, I considered the possibility that my study was under-powered and simply required more participants. However, the statistics provide no indication of this (e.g., for completion time I have an F-ratio of less than one, with a very small η^2). While collisions could be considered a

*Table 4.7. Mean ranks and chi-square values for perceptual effects for no priming. *p<.05*

	unsafe	middle	safe	$\chi^2(2)$
speed	2.1	2.2	1.7	3.6
weight*	1.4	2	2.7	17.4
steering*	1.5	2.4	2.1	11
durability*	1.6	2.2	2.2	7.6
safety*	1.6	2.4	2.0	9.8
responsive	2.1	2.2	1.8	2.4

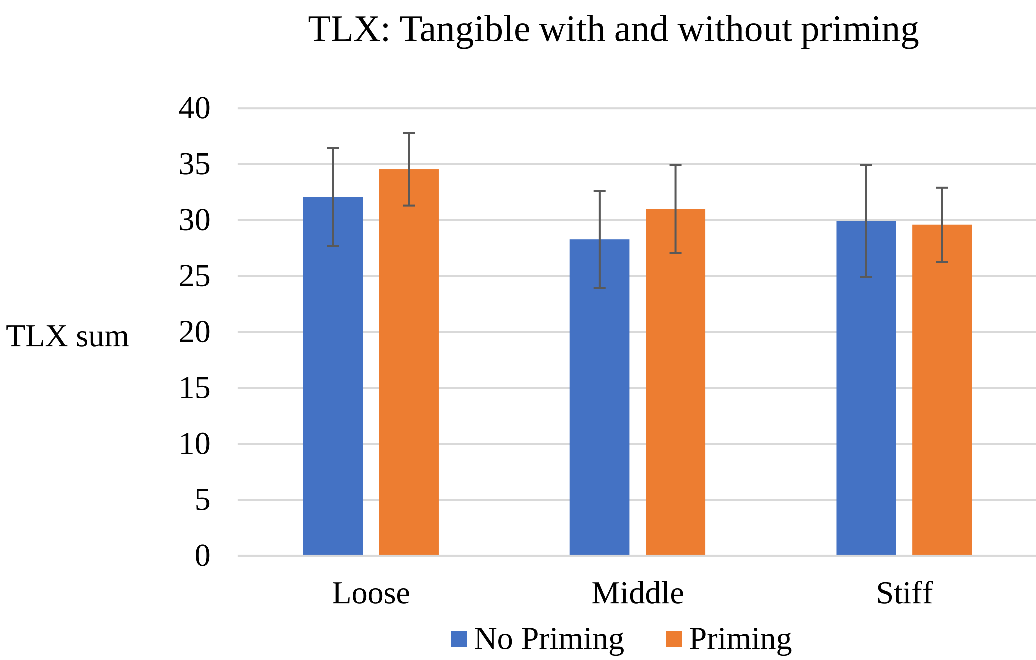

Figure 4.13. TLX results of the Tangible Priming condition (No Priming) with the Descriptive Tangible Priming condition. Just tangible priming did not find a difference in TLX score, but our results are inconclusive. Error bars are 95% CI.

trend with a medium effect (p=.07, η^2=.14), the effect was opposite of the prior study (stiffer joystick had more collisions), and the actual differences observed were much smaller (on average 1.9 collisions versus 4.8, Figure 9), suggesting that the three joystick stiffness levels were similar in this unprimed case.

The lack of a workload difference in my no priming study also supports my original conclusion that the difference in workload was due to my priming method. However, I saw an F-ratio of 2.2 with a medium effect size (η^2=.11) and a roughly similar trend of lower workload as stiffness increased (Figure 4.13). This suggests a small effect may be seen with more participants. This leads me to believe that usability plays at least a small part in my tangible priming workload change, with perhaps the remainder played by priming. I emphasize that this result is inconclusive for workload.

Thus, in the tangible priming case, operators felt a difference in the driving feel (change in workload) and did drive differently (a significant change in number of collisions). In the no-priming case, operators did not feel the workload was different, and did not drive differently. If change in joystick stiffness really was a major usability factor and usability created the reduction in collisions I originally observed, I would expect it to be reflected in both studies. My observations did not see similar changes in my no priming study.

Thus, I conclude that it was likely my priming method primarily (tangible priming), and not the usability of the device, that resulted in better driving behaviour. This is a striking result as I was able to change driving behaviour simply by inferring robot capability to robots through joystick stiffness.

4.6.4 Discussion – Tangible Priming and Collisions

In my original study, I found a decrease in collisions for the stiff (safe) priming when applying tangible priming. If I take the results of the no priming experiment into effect and compare them to tangible priming, it appears that tangible priming may have actually *increased* the number of collisions (Figure 4.14). While this may suggest we should not use priming at all, I see both experimental design choices and analytical issues that limit this conclusion.

In my tangible priming study, I told participants that changes to the robot may be felt through the joystick, and that I was investigating what they thought of each. Thus, participants had to test the tangible feel, interpret what the changes mean, and continually re-evaluate these

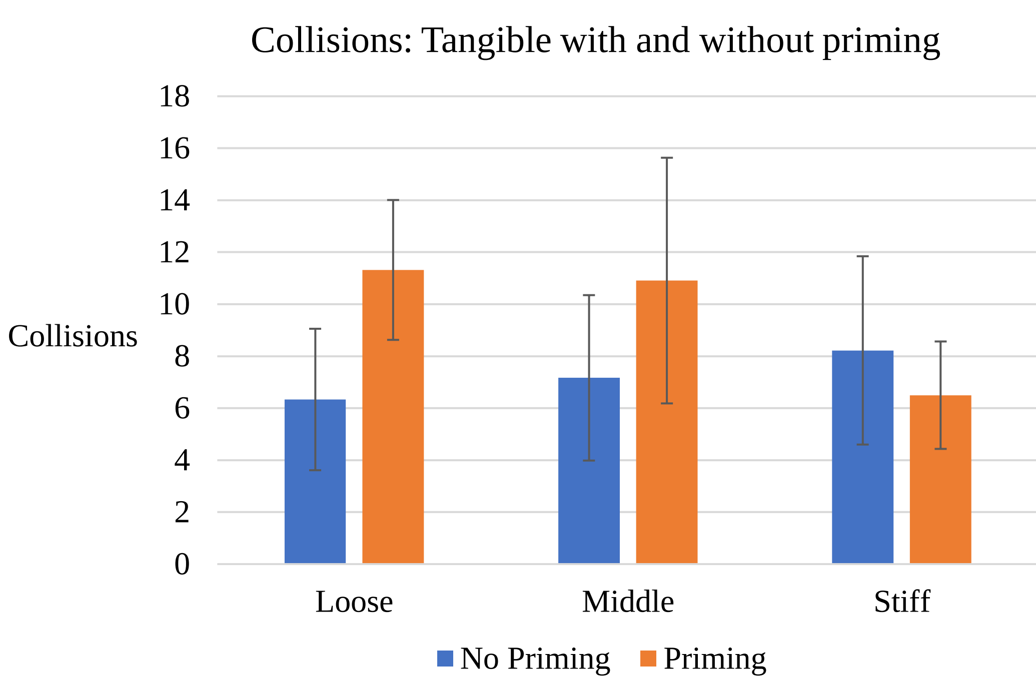

Figure 4.14. Collision results of the no priming (just joystick stiffness change) with the Tangible Priming condition. The no priming condition did not find a difference in collisions and has a trend moving in the opposite direction of the tangible case. Error bars are 95% CI.

thoughts while operating the robot and constructing their own opinions of driving the robot. The no priming study, in contrast, has operators only need to interpret what they thought of the different joystick feel; there was no need to interpret vague instructions on what may be changing in the robot. In this sense, the experimental design itself, or what I asked of participants, created an extra cognitive load for the operators, perhaps impacting performance.

In terms of analytics, I wanted to compare the two studies statistically: they act as one set of subjects with priming method as a between-subjects factor. However, these subjects were not recruited at the same time of year, and are not from the same set of potential participants (a full academic year afterwards). To find if there were differences between the cohorts of the two studies, I compared the demographics with a one-way ANOVA, but did not find any statistical differences, however ($p > .05$). Seeing the cohorts were similar, I wished to see if the number of collisions were different between the two experiments. Thus, I performed a post-hoc repeated measures ANOVA with priming as a between-subjects factor. Priming was not found to have a main effect on time ($F_{1,38}=.1$, $p=.4$, $\eta^2=.003$), collisions ($F_{1,38}=1.4$, $p=.25$, $\eta^2=.04$), or TLX scores ($F_{1,38}=.4$, $p=.5$, $\eta^2=.01$). While it may appear that the use of priming increased collisions in my data, both statistics and my study design make it difficult for me to support this claim.

4.6.5 Follow-up Study Discussion – Why did perceptions of the robot change?

Despite finding no differences on driving ability or perception of workload, I did find that operators rated the robots' perceived capabilities differently based on the joystick stiffness (Table 4.6). This is a very curious result, particularly given how I explicitly told participants (repeatedly) that they were driving the same robot multiple times.

The relative ranks for perceived speed, durability, and steering shared some similarities between tangible and no priming (Table 4.8). The shape of the results, the relative rankings, was shared for speed. The loose joystick was also seen as the least safe in both conditions for steering and durability. In each case, overall safety aligns with different perceptual measures. This implies that no single perceived quality decides perceived safety: feelings of speed, feelings of control (steering), consequence (durability), and performance (collisions) together can produce perceptions of safety. While my NASA TLX results do not necessarily refute this statement, I note that the TLX scale measures more than simply cognitive load, complicating this interpretation.

We considered several possible explanations for why operator perceptions of the robot changed in the no-priming case, even though I told them they were driving the same robot. One possibility is that some participants may simply be reporting on *the entire system* (robot, computer, joystick, etc.) and not just the robot; in this case, the joystick did change, which was then interpreted as a change to the robotic system as a whole. Another possibility is that my study design, with me repeatedly asking the same descriptive questions after each case, may have made participants feel pressured to compare the conditions and find differences. However, this does not explain why the differences reported were consistent across participants: if there was no consistent effect of stiffness on perception I should have seen noisy results, not statistical significance on multiple measures. Conceivably then, despite knowing that the robot did not change, the participants reported on how the joystick setting made the robot *feel*; perhaps this constant tangible reminder still constituted a priming effect as described in my tangible priming design section. This highlights the potential complexity

Table 4.8. The perceptual rankings of the tangible and no priming studies have been reproduced here for comparison. Note that for No Priming, the speed measures were not found to be significant, though all otherratings here are.

	Tangible Priming			No Priming		
	loose	middle	stiff	loose	middle	stiff
Speed	2	2.4	1.7	2.1	2.2	1.7
Steering	1.6	2.2	2.2	1.5	2.4	2.1
Durability	1.7	2	2.3	1.6	2.2	2.2
Safety	1.7	2	2.3	1.6	2.4	2.0

of considering priming or trying to remove its effects, with many different parts of the user and system interacting and changing how a robot is perceived and used. I reflect on this complexity more in the next section.

4.7 Reflection on Priming: Unavoidable and Complex

To draw insight about priming in teleoperation as a whole, I compare the results of all three studies together below.

4.7.1 How to (Not) Avoid Priming Effects

In the previous section, I briefly considered that priming may be possible even when my operators were explicitly told that only the joystick stiffness was changing. This relates to prior work that demonstrates how priming can be effective, even when participants are aware of the priming attempt (Doyen et al. 2012; Cheesman and Merikle 1984). Considering this, I note the similar results between the no priming and the tangible priming study discussed previously, as well as how no priming and descriptive priming also shared similar ranks in perceived weight and safety for all three conditions. It is statistically unlikely that these similarities between the two priming and no priming methods are just by chance. Thus, I have to consider that the joystick stiffness may still have served to prime my participants, encouraging them to draw from prior experiences and understanding to shape their perceptions.

This highlights the complexity of considering priming with interface design and creates a problem for my prior analysis. If my no priming study can still be considered priming despite my attempts to counteract it, then even if I saw changes in driving performance, I could not conclude whether it was due to usability or priming: I could not fully remove the priming element to isolate and measure the base usability effects of joystick stiffness. This line of reasoning highlights that it may be impossible to avoid priming, or to separate other effects (e.g., usability of an interface) completely from how the design and presentation will prime users. Thus, I argue that it is important for the field to more explicitly consider the priming effects of their interface design, even when not a major intended component.

4.7.2 Complexities of Priming Methods and Effects

While reflecting on my three studies, I observed that many small considerations in priming design may make a big impact on the end effects. Even small variations in how I presented a priming method may explain differences between my tangible and no priming cases. I caution that because of this, even wording and information granularity in introductory or marketing materials may affect perceptions and possibly operator behaviour.

Comparing the priming designs of all three studies together, I realized that my initial tangible priming shared similarities with descriptive priming: I *described* the connection between the (supposedly different) robot and the different tangible feelings in the joystick by implying the robot changes could be felt (somehow) via the joystick. In other words, I linked feelings of control to different robot characteristics in an indirect way, and so when participants felt differences in the joystick they possibly tried to reason about what characteristics in the robot may make the control method seem stiffer. This is another potential explanation to why perceptual changes were different between no priming and tangible priming, discussed above, and further demonstrates how small differences can make a difference in priming. It would be illuminating to compare tangible priming to a version where the explanations of the tangible differences are more specific than I was.

In my no priming condition, participants were instructed that the tangible sensations had no meaning. This may imply there it was not rational thought that changed perceptions, but a more subconscious feel of the controls. That is, it seems that the interface *feel* impacted participant perceptions of the robot capability despite them knowing that the robots were the same. This is related to the aforementioned vagueness of my up-front explanation during tangible priming; guiding the interpretation like I did in descriptive priming may be able to better target priming's effects. This suggests that guiding the interpretation of potential priming stimuli is be important, though doing so on a subconscious level may be difficult and hard to measure.

Operators' explicit knowledge that they were driving the same robot may stop some qualities of a robot from being primed. Tangible and descriptive priming both changed an operator's perception of robot speed, but this was not observed in the no priming experiment. Note that operators were either told some robots were faster (in descriptive priming) or free to imagine different speeds (in tangible priming). I suggested priming is impossible to avoid, but it is perhaps avoidable for certain robot properties, depending on the presentation of how the robot should be perceived. While my work was not conclusive in this regard, it is an important direction as perception of speed can interact with driving safety and behaviour (Recarte and

Nunes 1996; Fuller 2005).

Further adding complexity, the combination of tangible and descriptive approach in tangible priming may have leveraged multiple priming methods. This may have "boosted" the priming signal to be stronger, which has been shown to be an important factor determining the effects of priming (Cheesman and Merikle 1984). Thus, priming method variants as well as combinations with other methods may all result in unexpected effects.

4.7.3 Priming Operator Perceptions with or without Changing Robot Capability

When changing robot capability, I found that reducing a robot's ability to move quickly could increase safety without having a proportionately large trade-off in completion time. Interestingly, the operator's perception of their own performance also increased, along with a decrease in the effort they perceived they used. Simply increasing robot ability may not be a useful strategy in some cases; both the user's needs and the demands of the task should be taken into account when designing interaction. Further, it may be that a user feeling like they are in full control of a robot will perform better. These effects mirror those seen in video games, where players may choose virtual vehicles with lower top speed as a trade-off for improved driving in other ways, such as ability to control the vehicle.

For tangible and descriptive priming techniques, I did not change the robot's abilities, but still could change an operator's perception of the robot's ability to be driven safely. It is interesting that changes in perceived robot capability lead to similar performance differences as driving profile priming where I actually adjusted the robot's abilities – both encouraged safer driving and lowered cognitive load. This implies that change in robot ability may have a similar impact as changes in perception of the robot in terms of influence on operator performance. My evidence does not completely support this, as I did not see safety differences with descriptive priming, but highlights both reducing robot capability and some priming methods as an exciting potential method to improve teleoperation experience and performance.

4.7.4 Why Consider Priming in Teleoperation Design?

The deeper I try to push my interpretation of my results, the more complex the picture of priming methods, their design, and their effects, becomes. Not only does this complicate my interpretation of the negative results in my no priming condition, but it can also explain some of my confusing findings, such as why people reported perceiving the robots differently *even though I told them they were the same*. Similarly, my tangible and descriptive priming methods shaped perceptions of the robot and convinced participants that the robots were different, despite driving the identical robots for upwards of 30 minutes; given that, in all cases, the robot was identical and did not change, one would reasonably expect participants to drive approximately equivalently, to rate the robots similarly, and to realize that they are the same or very similar. That I found differences in all the above measures reflects the potential of priming, and its potential in teleoperation. Further, perceptual changes alone could useful immediately for teleoperation, as perceptions of a product and its quality can heavily influence product success and use patterns (Bargh, Chen, and Burrows 1996; Lindgaard et al. 2006; Mitra and Golder 2006; J. E. Young et al. 2009). Designing robots to *feel* safer or more usable can affect use, adoption, and popularity.

4.8 Limitations

While my priming methods were successful in changing participant perceptions of the robot and teleoperation experience, I only found teleoperation performance changes with the tangible and driving profile methods. I discussed potential reasons for those results above, but I note my quantitative measures in all three studies were not exhaustive; exploring other performance metrics (e.g. average robot velocity), or trying to better measure a teleoperator's driving skill will help me better understand the limits and potential of priming on teleoperation performance. Regardless, as perceptions of technology can affect user experiences and influence adoption and acceptance of technology (Bargh, Chen, and Burrows 1996; Lindgaard et al. 2006; Mitra and Golder 2006; J. E. Young et al. 2009), even without performance differences, priming can be an important tool for roboticists in shaping how their robots are perceived and accepted.

This work assumes that people respond in the same way to each priming stimulus. However, it could be that different personalities may be more prone to risk taking, as suggested in transportation research (Jonah, Thiessen, and Au-Yeung 2001). Further, people have varied driving experiences, both in terms of frequency and types of vehicles driven, and this will change what experiences a priming method can help a person recall. In my results, my *safe* condition primed *safe* behaviour, while some previous research suggests that the inverse may be true; for example, adding safety features to cars may result in less safe driving (Jonah, Thiessen, and Au-Yeung 2001). In teleoperation, a fast robot may encourage safer driving behaviour from a cautious person, or a thrill-seeking operator may get excited and try push the robot to its limits. I note that the science surrounding priming still has conflicting results (Doyen et al. 2012), thus I recommend further inquiry into priming and teleoperation, considering a participant's risk-tolerance.

My scenario also limits the generalizability of my results. The obstacle course was designed to imitate a very crowded office or conference venue and make teleoperation difficult. However, environments with dynamic obstacles (such as people in a busy subway station), or wider spaces such as many museums will change the teleoperation experience. As I noted earlier that research suggests that context is important for priming effects, investigating context for teleoperation and priming is an important consideration. I further noted this as a potential explanation for why my acceleration-limited robot did not perform well – its abilities may simply have not been suited to the task specifications. Priming my operators may have made them believe that one robot was more suited to the crowded obstacle courses, which explain perceptual or performance differences I observed.

4.9 Future Work

My results serve as a base to build from for future priming-based teleoperation interfaces. Even my two priming method labels – descriptive and tangible – are general and can be explored much further and much more deeply. For example, descriptive research may look at priming with actual demonstrations of robot behaviour (using acting to prime the danger or ease of teleoperation), different robot morphologies, or different robot sounds. Similarly, additional tangible methods could control force feedback effects such as adding shake to simulate rough terrain or a powerful motor. Exploring each technique in depth and starting to explore a broader range of priming techniques, is important for understanding the nuances of how priming can affect teleoperation.

I should also explore priming beyond portraying the robot as more or less safe. For example, sound could be used to represent environmental danger in real time, or I could explore whether the enjoyment of teleoperating the robot could be primed. This is a new avenue to consider for teleoperation robot and interface design, and it leads to a broad range of future work.

Priming effects are often studied in the short term, such as my work in this paper. Long term effects of priming are less studied, and thus should be studied in the context of teleoperation; prior work suggests priming may last for hours or even months, even if new experiences contradict the priming (Sloman et al. 1988; Lindgaard et al. 2006; Becker et al. 1997). Perhaps short-term priming effects, especially when operators are first learning to drive a robot, may influence the development of safe long-term habits, but this must be formally studied. Such research would benefit both the psychology and teleoperation communities.

The mystery of how my no priming condition resulted in changes in perception of the identical robots people drove is also an important avenue to understanding priming. Part of the difficulty in pursuing this reason is my use of participant-volunteered responses; while I believe qualitative feedback is very important to understanding participant reactions to priming stimuli, it is inherently interpreted first by the participants themselves which makes it difficult to understand true causal relationships in priming. I recommend future studies couple measures like I used with other, perhaps new techniques, to measure and understand a person's internal thoughts, dialogue, and even subconscious processes when being primed.

4.10 Conclusions

In this chapter, I introduced some methods video games use to influence how players think about and use their in-game characters. In particular, I looked at how games, particularly racing games, may shape expectations of what a character can do, and how this may affect how players play the game, regardless of if that difference is true or not. I combined this inspiration with the priming – shaping expectations to change behavior and perception.

I demonstrated how different priming methods are able to impact a teleoperator's perception of a robot, their experience teleoperating it, and their teleoperation performance. I took priming, which has been used and studied extensively in psychology, and presented it as a concrete and practical tool to be used by robot teleoperation designers.

I explored multiple priming approaches: priming with driving profile, tangible, or descriptive methods. My results demonstrate how priming can be successfully used to change operators' perceptions of a robot's speed, safety, power, and weight, and more, even when operators believe no priming is taking place. I also demonstrate priming's ability to affect an operator's driving behaviour and improve safety. These changes can occur without ever changing the robot, its programming or behaviour, or on-screen interface.

Changes to actual robot ability as a form of priming does indeed change performance, as expected. However, I found that changes to speed and acceleration did not produce the results I originally expected. Performance changes affected operator perceptions of their own performance and effort, with slower robots yielding higher ratings for performance, lower effort scores, while reducing collisions. Taken with my other priming methods, I found that changes in actual robot performance can create similar performance improvements as my tangible method, but not my descriptive priming method. This raises more questions about the links between perception and performance (of both robots and operators) in teleoperation.

Interface and robot design continue as challenges to improve teleoperation, as both are important for usability and user experience. While my series of studies and meta-analysis opened even more questions, I believe that this work on priming teleoperators provides human-robot interaction designers with an additional tool to further shape teleoperation performance and user experience.

In addition to demonstrating the effects of priming, I showed how non-functional design changes and presentation can impact the operator's behaviour and experience. Video games often concentrate on aesthetics and presentation to create atmosphere or suggest ways to play the game. My results here suggest that this may be an important area of future work for teleoperation in general.

CHAPTER FIVE:

SOCIAL INTERFACES IN TELEOPERATION

In the last chapter, I demonstrated how video game inspired techniques can influence operators to perceive and operate robots differently. Another way video games can influence players to act differently is to use social techniques, such as having computer-controlled characters with different personalities present different options, or have a social agent interface, such as an on-screen face that presents data using social signals. The core idea is that people are inherently social, and process social feedback quickly and even subconsciously (Lee and Nass 2010; J. E. Young 2010). I explore how the presentation of data with a game-inspired social interface may affect teleoperation.

Social human-robot interaction has found that robots can leverage this social nature of people (J. E. Young 2010) to diffuse arguments (Jung, Martelaro, and Hinds 2015), build empathy with a person (Seo et al. 2015a). Video games similarly use social communication to help convey state, increase awareness of the environment, or influence player behaviour. However, this use of social communication has yet to be explored to teleoperation. I explore this intersection of video games, teleoperation, and social human-robot interaction to see how social techniques can affect operator behaviour and experience.

Specifically, I explore if an on-screen agent that reacts to a teleoperator's driving performance can influence the teleoperator and their driving; for example, the agent could show fear during poor driving to perhaps influence the operator to slow down. My design concept is to have the agent feel like a virtual passenger or companion to create an emotional response in the operator (e.g., to feel bad for the agent) and ultimately can shape behavior (e.g., to slow down to calm the agent). This is inspired by characters in games that frequently react to in-game events and are given dialogue or emotional interactions with the player as part of the game, building empathy and eliciting emotional reactions during gameplay.

I evaluated this concept by designing and implementing two proof-of-concept agent personas that react in different ways to operator driving. By conducting an initial proof-of-concept study comparing my agents to a base case, I was able to observe the impact of my agent personas on operator experience, perception of the robot, and driving behavior. My results demonstrated that emotional on-screen interactive agents, like their video game inspirations, can alter teleoperator emotion and may even affect teleoperation behaviour. My results highlight potential for more targeted, ongoing work in applying social techniques and strategies from video games to teleoperation interfaces.

Parts of this chapter have been taken in part or in full from the following publication:

Daniel J. Rea, James E. Young, "Backseat Teleoperator: affective feedback with on-screen agents to influence teleoperation." The ACM/IEEE International Conference on Human-Robot Interaction (HRI'19). ACM/IEEE. 2019. (24% acceptance rate)

5.1 Social Interfaces for Teleoperation

I present an approach to improving teleoperator performance that aims to use video game-inspired social feedback to impact an operator's mental state, with the ultimate goal of trying to shape how they drive the robot. Specifically, I add an interactive agent to a simple teleoperation interface, like a virtual passenger, which reacts to operator driving using emotional feedback. Ideally, the operator may feel empathy and compensate by altering their driving (Figure 5.1). For example, if the agent acts scared following a collision, the operator may feel empathy and automatically drive more safely to console the agent. This effect, of a person witnessing an emotion and, in response, changing their behavior or feeling an emotion themselves, is well-established in other fields (e.g. Schoenewolf 1990; Hatfield, Cacioppo, and Rapson 1994; Barsade 2002). In this chapter, I present and explore a proof of concept using this technique to shape teleoperator experience with the intent to change operation behavior.

Social human-robot interaction has found that robots can use human- or animal-like social communication techniques when working with people in an attempt to improve and simplify communication with them (J. E. Young et al. 2009). For example, autonomous robots co-located with people can use techniques such as expressive movement (Sharma et al. 2013a), gaze (Breazeal et al. 2005), or even animal-like tail movements to convey robotic state or intention (Ashish Singh and Young 2013). However, apart from social teleoperation (where a robot is a proxy for two remote people interacting), there has been little work done that explores how a teleoperated robot can similarly use social techniques in the interface to support their operators. As such, I present this work as a proof of concept, where a teleoperated robot aims to use techniques from social HRI to impact the teleoperator.

Social interfaces are also used in video games. They can be used to build rapport with the non-player characters and the player (e.g., Mass Effect 2, BioWare, 2010; Red Dead Redemption, Rockstar Games, 2010), can talk and bring attention to important ideas and objects in the world (e.g., The Legend of Zelda: Ocarina of Time, Nintendo, 1998), or even build empathy between the player character and the player themselves (e.g., Doom, id Software, 1993). This can affect the player's decision making, especially as many of these social interfaces are there throughout the game, build relationships with the player, and react in a personal manner despite being virtual (Figure 5.2). The Doom example in particular serves as my inspiration, which tries to leverage an on-screen *social visualization* of the teleoperated robot's state (Figure 5.3).

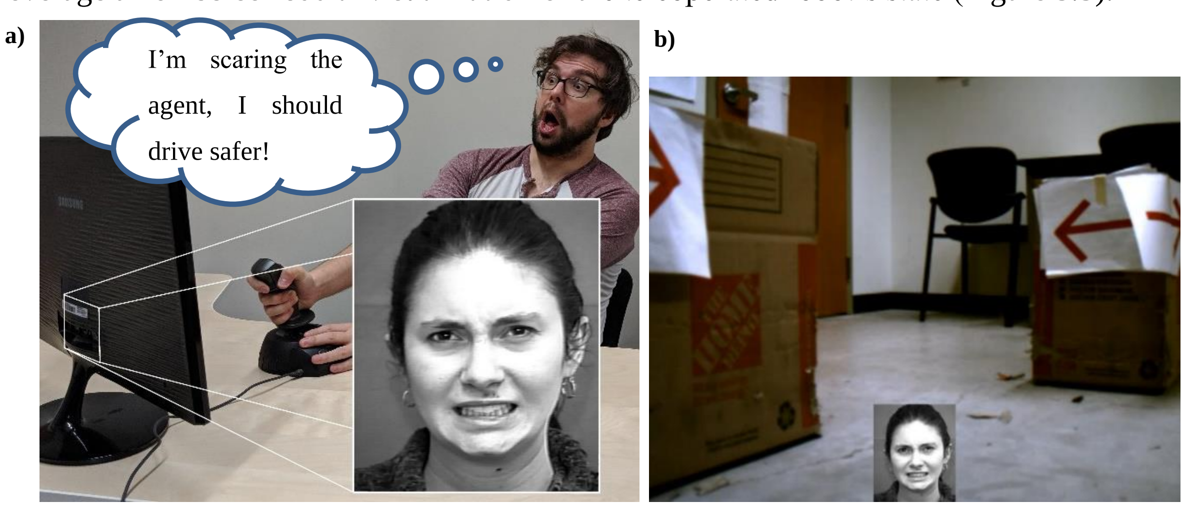

Figure 5.1a) An on-screen "virtual passenger" agent reacts to poor driving by exhibiting anxiety, with the intention of impacting the teleoperation experience (dramatic re-enactment) b) the interface displayed during robot teleoperation.

We designed and implemented a virtual passenger for the teleoperation interface which reacts to the operator's driving (e.g., average speed, collisions) in real time by displaying an emotion (Figure 5.1), with the goal of shaping the teleoperator experience and perhaps ultimately their driving behavior. For example, if the agent reacts with a positive emotion, such as smiling, the operator may similarly become more positive which may reinforce the current driving style. Conversely, a negative agent emotion, such as anxiety or fear, might discourage the current driving behavior. To explore this approach, I designed two agent variants, each using a different affect feedback model, and conducted an initial study to investigate how these agents may impact the teleoperation experience and operator's driving.

My results indicate that affective-feedback passenger agents can create emotional change in teleoperators. However, I found no compelling evidence that they changed driving behavior in this case; my analysis highlights limitations and avenues for improving both the agent and study design that will be useful for follow-up work. Overall, my work serves as a proof of concept of using affective feedback-based interfaces in teleoperation, and of using social

Figure 5.2. In-game decisions in Mass Effect 2 (BioWare, 2010) can be tense, with your computer controlled companions giving suggestions, sometimes strongly, about what the player should do. In this example, the players companion ("Wrex") draws their weapon on the player because they find the player's intended actions to be immoral (blowing up the base). The player can interact socially to disarm the situation. Further, these companions follow you throughout the game, building a strong relationship, potentially lending weight to their arguments, even though they are just virtual characters. In this case, two AI companions are ideologically opposed, and the player may even shoot one in the argument.

Figure 5.3. A screen shot from Doom (id Software, 1993). The player avatar's face is displayed on the bottom, middle of the screen. It reacts to the surroundings (damage from enemies, surprise attacks, excitement for new weapons), building empathy for the avatar. We use this interface as inspiration in our design.

interaction techniques to support operators in general, which I envision will be an important research topic for teleoperation moving forward.

5.2 Social Psychology and Teleoperation

In this thesis, I have demonstrated how different components of human psychology play roles in affecting teleoperation performance and experience. In particular, I have shown how cognitive factors and changing expectations are important – subtle cues can leverage psychology to improve the operator's behavior (Daniel J Rea, Seo, et al. 2017; Daniel J Rea and Young 2018; Hacinecipoglu, Konukseven, and Koku 2013). I further extend this line of inquiry to the use of social psychology in teleoperation interfaces.

Research in traffic psychology has shown that a driver's psychological state can change their driving behavior (J.A. Groeger and Rothengatter 1998; John A. Groeger 2002). These changes may be due to the perception of the vehicle itself (Eyssartier, Meineri, and Gueguen 2017), the surrounding environment (Michon 1985) or even the physical controls of the vehicle (Blommer et al. 2017; McIlroy, Stanton, and Godwin 2017). Importantly for this work, the driver's mood may be a factor in driving safety (Precht, Keinath, and Krems 2017). This body of work demonstrates that a driver's mental state or emotions can influence how they drive. I build upon this base of traffic psychology and investigate if I can use affective feedback to change an operator's emotions with social interfaces, and further investigate if this changes their driving behavior.

The use of social behaviors and strategies follows an established tradition in social robotics, and human-computer interaction in general (Lee and Nass 2010). For autonomous robots, the use of social behaviors has been shown to influence group communication dynamics (Jung, Martelaro, and Hinds 2015; D. Sakamoto and Ono 2006), dissuade people from performing actions (Briggs and Scheutz 2014), encourage lying to authorities (P. H. Kahn et al. 2015), or change how people talk (Brandstetter et al. 2017). I see these examples as demonstrating an opportunity to have robots use social phenomena to change and affect interactions and people's behaviors with them (Postnikoff and Goldberg 2018; Sanoubari et al. 2019). Social behaviors have further been used to communicate robot state (e.g., A Singh and Young 2012; Sharma et al. 2013; Breazeal et al. 2005; Young, Sharlin, and Igarashi 2013). I extend and combine these strategies by using social behaviors *in the teleoperation interface* to communicate state and simultaneously influence the teleoperator themselves.

Social feedback in vehicle driving situations has been shown to be beneficial, such as in car interfaces (Leshed et al. 2008). However, the design of such interfaces is non-trivial, and may be distracting (Srinivasan and Jovanis 1997) and increase cognitive load (Blanco et al. 2006; Drury, Scholtz, and Yanco 2003). My design aims to explore emotional displays as a social feedback mechanism, while also exploring how the effects may change teleoperation behaviors.

Social signals and teleoperation are often studied together in the context of telepresence. Telepresence research tries to design robots and robot interfaces that are used by one person to control a robot and interact with another person socially, where the robot is a proxy (e.g., Rueben et al. 2017; Kristoffersson, Coradeschi, and Loutfi 2013; Tsui et al. 2013; Tanaka et al. 2016). My work contributes to teleoperation by using social feedback mechanisms in cases where there is no human on the remote end: the social interaction is between the operator and the teleoperation interface.

5.3 Design: Interactive Teleoperation Agents with Affective Feedback

As a proof-of-concept for using affective feedback in teleoperation interfaces, I designed two interactive agent personalities to influence an operator's mental state and potentially robot driving behavior. The agents monitor teleoperation performance in real time, and based on how well the operator is driving, the agents change their facial expression. To explore this space, I designed two different agents, each with a specific affective feedback and reaction strategy. I note that there is a rich potential for future work in applying more complex and thorough psychological frameworks to agent design; my goal here was rather as an exploratory proof-of-concept with agent designs that follow a simple model.

My design was heavily inspired by the video game *DOOM* (id Software, 1993), where the face of the player's avatar was displayed at the bottom of the screen and reacted emotionally to the avatar's state and events in the environment (Figure 5.3, Figure 5.4).

5.3.1 Design Strategy: affective feedback

Our approach to influencing an operator is to leverage affective feedback by showing them an emotional reaction to their driving. Previous work has found that when a person sees someone experience an emotion, the viewer may experience a similar emotion (becoming happy when

Figure 5.4. A partial screen shot from Doom (id Software, 1993), and sample faces of the player character's reactions. Our samples show how the player character can appear to be grimacing at something to the left (left), being suspicious of something to the left (middle), or reflecting character state by becoming bloodier as the player character takes damage (right). The face updates frequently throughout play.

someone around you is happy), often an automatic or reflexive response (Hatfield, Cacioppo, and Rapson 1994). Alternatively, if the operator develops empathy for the agent they may react by trying to support the agent (Bartneck et al. 2007; Seo et al. 2015a).

My goal is to use affective feedback to induce an emotional response in the operator. I do this for the purpose of shaping driving behavior and teleoperation experience. My exploration concept is that positive emotions will influence behaviors via positive and negative reinforcement: the happy face may make the operator feel happy as well, providing positive reinforcement for the driving behavior at that moment. Conversely, I expect my affective feedback will create negat ive emotions in the operator if the agent reacts negatively. I expect this to provide negative reinforcement and dissuade the operator from taking similar actions in the future. With these two ideas in mind, I designed respective interactive agent personas with different affective feedback strategies: an "anxious" and a "daredevil" agent.

5.3.2 Personas for Affective Feedback

Both personas are based on the same principle of trying to encourage certain behaviors with positive emotions and discourage others with negative emotions. Specifically, my agents encourage or dissuade behaviors based on teleoperation danger, such as collisions with obstacles, or driving too quickly. Thus, the reactions act as a social interface that conveys safety information to the operator.

Anxious persona: if an operator drives more dangerously, the agent would become more upset or frightened. Conversely, if the operator drove safely, the agent would become happier. This was to encourage safe driving with happy reactions and dissuade less safe driving.

Daredevil persona: the agent displays an increasingly bored and contemptful face if the operator drives safely but becomes excited if driven dangerously. I expected this persona to promote dangerous driving by providing positive affective feedback when the operator drives dangerously. Further, the negative reactions to safe driving may discourage safe behavior. This was designed to explore if a poorly designed persona could possibly promote dangerous behavior.

The daredevil and anxious personas both build on the same approach of leveraging social feedback to change teleoperation behavior, with the different personas helping to explore my strategy.

5.3.3 Measuring Teleoperation Safety

This initial proof of concept uses collisions per minute and robot velocity as coarse measures of driving safety. Collisions are a direct sign of mistakes during operation. Velocity is a measure of safety as, in general, driving very quickly is more dangerous: faster speeds give operators less time to react and not collide with people, expensive equipment, or tumble over a ledge. I acknowledge that a very skilled driver may be able to drive quickly without causing collisions, but they are still subject to these increasing constraints to reaction time and may still make a real (or in my case, virtual) passenger nervous. I concede that my choice of these two measures is a limited representation of safety – reckless acceleration, near misses, and other factors may all contribute to long-term safety. It serves, however, as a sufficient and consistent mechanism for my initial exploration.

5.3.4 Design Implementation

I designed my interactive agent to be easily visible but to not be too distracting. This was done by placing the agent on-screen, overlapping the teleoperation video in a salient location while not covering up a typically important area (Figure 5.1b). Further, to provide an illusion of activity for the agent and draw attention (Franconeri and Simons 2005), I had the agent update its expression twice per second. That is, even when the reaction did not change due to the operator's driving safety, the face would have small changes to support the illusion of constantly reacting.

Calculating Safety

In order to define how my agents reacted to teleoperation, I had to define a number of collisions per minute and speeds were considered unsafe or safe. I ran pilot experiments to calibrate this, specifically tuning the change in velocity or collisions per minute needed to change the reactions of my personas. My goal was to find thresholds such that the agents provided noticeable visual and emotional feedback for both the operator's initial driving, and after any changes they may make to their driving in response to the social feedback. Thus, my thresholds are specific to my environment, and are admittedly limited and ad-hoc.

I calculated an independent safety rating for both collisions per minute and driving speed, resulting in a value that ranged from most safe to least safe. For collisions, I maintained a running "collisions per minute" total, which summed collisions occurring in the last minute. I used a linear weighted sum to make the agent's changing reaction smoother as older collisions became less relevant: each collision was weighted by how much of a minute had passed since the collision occurred. For example, a collision that was 30 seconds old would contribute to the safety rating as half a collision. Collisions were measured automatically by combining data from the robot's inertial measurement unit and the joystick used to drive the robot.

Velocity-based safety was calculated based on the average velocity over the last minute. I defined "not safe" driving to be anything over a threshold speed (25% of robot max speed). Excess velocity after this threshold was then used to determine the safety rating. As discussed earlier, I did not want to react to maximum speed driving with no collisions as completely unsafe. Thus, max velocity safe driving (no collisions) would only progress the personas to a middle safety state (Figure 5.5, neutral).

The final safety rating was calculated as the least safe of the two measures, collisions per minute and velocity, recalculated each frame.

Selecting a Reaction

My interface maps the safety rating, ranging from a minimum safety rating to a maximum rating, to a reaction (Figure 5.5). I first ordered the persona's expressions from least safe to safe. I then use my calculated safety rating as an index in between these expressions; for example, a safety rating of 50% of the maximum safety rating will pick a neutral expression (half way between not safe and very safe expressions). A safety rating of 75% would pick a slightly smiling face, in the case of the anxious person (Figure 5.5, top).

My expressions are taken from validated video data-set of people making pre-defined emotional reactions starting from a neutral expression (Lucey et al. 2010; Kanade, Cohn, and Tian 2000). The personas are formed by reversing the "not safe" emotion video to start from an emotion and end with a neutral expression. I can then transition to the "safe" emotion video by moving between the neutral expressions in both videos.

Thus, each expression in my dataset is a frame in this linked video – a video of an unsafe reaction, transitioning to a neutral reaction, transitioning to a safe reaction. My safety index is mapped to a frame in this video, which is displayed in my interface (e.g. Figure 5.1b). As the safety rating changes, I simply display new frames from the video, providing a smooth emotion transition. If the safety rating stays the same, a nearby, similar frame of video is used to show small movement in the agent, such as slightly moving the corners of their lips or eyes. This creates an illusion of activity and livelihood in the agent, and the movement may draw attention to the agent itself (Daniel J Rea, Seo, et al. 2017). My emotion video data is from the Extended Cohn-Kanade emotional face dataset (Lucey et al. 2010; Kanade, Cohn, and Tian 2000). In my anxious persona, I combined "fear" and "disgust" for unsafe reactions and used "happy" for safe reactions. In my daredevil persona I used "happy" for unsafe reactions and combined "contempt" and "disgust" for safe reactions.

5.4 Experiment: The Effects of Social Interfaces for Teleoperation

The goal of my experiment was to investigate the effects of my affective feedback interfaces

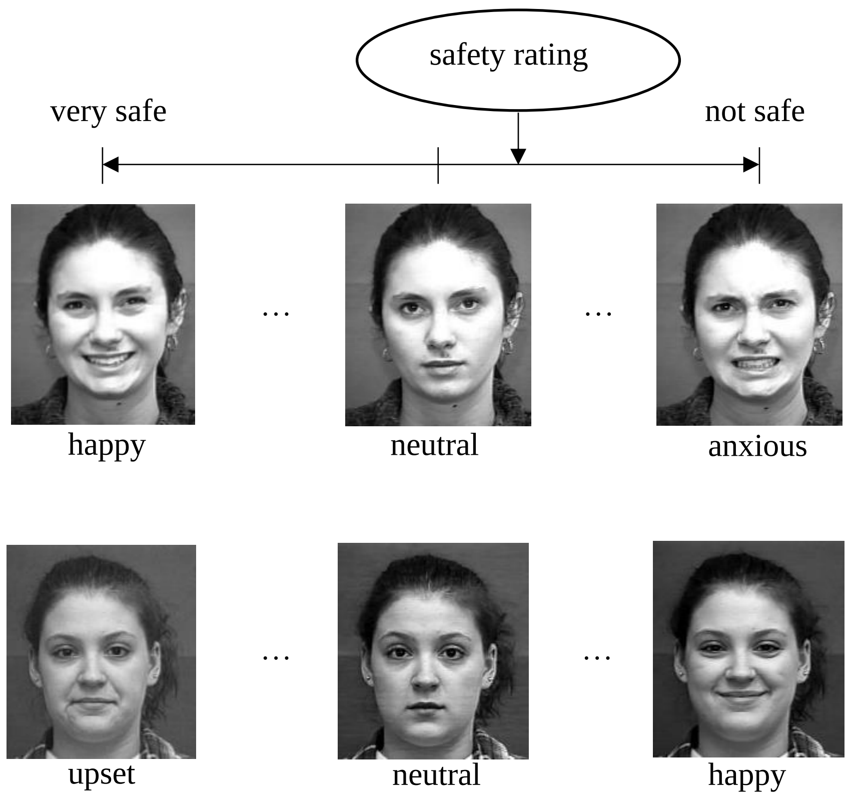

Figure 5.5. The range of expressions, mapped from very safe to not safe driving behavior. The real-time driving safety rating indexes into a collection of faces displaying emotion. Top row: anxious persona, bottom row: daredevil persona.

for teleoperation on the operator's perception of the robot and their driving behavior. To do so, I created a driving task: an obstacle course that would test a participant's ability to control the robot. Participants drove the robot through an obstacle course with the two interactive agents and a base case, and, in each trial, I measured their driving performance, their perceptions of the robot, and their driving using self-report measures.

5.4.1 Task

Participants were tasked with remotely driving a robot around an obstacle course. The course consisted of a grid of obstacles and a series of arrows that had to be followed, with each arrow indicating a 90 degree turn around a corner (heavily inspired from previous work Rea and Young 2018). Participants would drive 3 laps around the course, with the first lap being treated as a practice run. I instructed participants to drive as fast as they felt comfortable, while trying to avoid any collisions with obstacles along the course.

I designed three similar obstacle courses for the within-participants study; while the obstacles did not move between trials, the arrows leading them through the course did change. Each course had similar length and number of turns to maintain difficulty across conditions. Further, courses were designed to have a mix of straight sections and sequences of turns to test different driving scenarios.

5.4.2 Manipulations

I tested three conditions. The two interactive agents, *anxious* and *daredevil* personas, and a numeric-display base case. I struggled to develop a base case, as my first inclination was to simply have an interface with no feedback. However, this would compare two things: availability of driving feedback, and, emotional encoding. By including the numeric case, I can keep the feedback only without the affect. This base case displayed the same information encoded in my personas but had no social or inherent emotional element (Figure 5.6). Each persona started the condition showing the "very safe" reaction. This allowed me to test whether just the information alone could influence an operator's driving in comparison to the social encoding.

My experiment used a within-participant design; each participant used all interfaces: anxious

Figure 5.6. Our baseline interface simply displayed the safety information without social or emotional cues. The text reads: "collisions/min: 2.8 velocity: 68.6%"

persona, daredevil persona, and the baseline. Condition order was fully counterbalanced across participants, while course order was fixed for all participants.

5.4.3 Measures

Before the experiment, I administered a demographics questionnaire that recorded age and gender. I further inquired about any experience participants had for activities similar to robot teleoperation: experience playing video games, experience driving vehicles, experience with remote control robots (quad copters, RC cars, etc.), and participation in any other robot experiments.

In each condition I recorded the time it took to complete the task and number of collisions. During the experiment, I also logged robot velocity and the current safety rating of the participant's driving. The robot's movement data was recorded as a potential way to measure changes in operation.

To understand changes in self-reported workload, I administered the NASA TLX questionnaire (Hart and Staveland 1988). Further, I measured the operator's emotional state on a common two dimensional emotion model (valence and arousal, Posner et al. 2005), with the Self-

Assessment Manikin instrument (7-point variant, from -3 to +3, Morris 1995). To measure changes in perception of the robot's operation, we additionally asked participants to rate the robot's overall safety for driving, and informativeness of the safety indicator interface. The post-condition questions included free-form feedback areas for participants to give positive, negative, or other feedback that they felt was appropriate.

In my pilot studies, I initially noticed no effect, with participants noting they did not look at the face after the first few seconds. I thought there were two main factors for this. The first was that the original interface had the agent displayed off to the side, beside the video feed, which may have been difficult to see while concentrating on the main video. Also, I realized participants may not have felt compelled to pay attention as the task had nothing to do with the agent.

To encourage operators to pay attention to the safety information, I moved the interface to a less used portion of the bottom screen, as pictured earlier, and I created a distractor question about the information displayed. I ask operators to choose "the face shown most often while you drive," or "the average velocity you thought you were closest to most often." Then, I show a range of five faces used by the agent during the condition, spread from negative to positive emotions. For the baseline, five percentages of max velocity, spaced from 20% to 100% are shown. This question was not for analysis, but to encourage participants to pay attention.

After the experiment, I asked participants to rank each interface for preference. There were also optional short answer blocks for comments, similar to those described above in the post-condition questionnaire. Finally, I administered a questionnaire from prior work that measures susceptibility to emotional responses when exposed to different emotions, from (Doherty 1997). As my design was built on the premise that the agent's emotion could influence the operator's own emotion, I reasoned that the effect could vary wildly by how much a person could be influenced by an emotional display. This questionnaire measures potential susceptibility to different types of emotion (happiness, sadness, anger, etc.), and thus I used its ratings as covariates to control for variance in my participants.

5.4.4 Procedure

Participants were welcomed and told we will be exploring ways to convey safety information to robot operators and investigating how that may affect how safely they drive the robot. I explicitly told participants that I was using collision information and robot speed to gauge driving safety, and that this information would be displayed via a summary as a facial expression. I did not, however, state which expression correlates to what driving safety level. The consent form and demographics questionnaire were filled out at this point.

Both personas were introduced upfront and explained using a paper representation, with multiple expressions shown (similar to Figure 5.5). I additionally introduce the baseline system (Figure 5.6), and explained that the information it displays is the exact same information used by the system's algorithm to decide what face is shown (words like "system" and "algorithm" are used, emphasizing the mechanical nature of my interface, and not implying my agent is intelligent).

The participants were instructed that their task was to drive through the obstacle course as fast as they felt comfortable while trying to avoid all obstacles. After I explained the course instructions and controls, participants were given one lap to practice. Afterwards, they drove two laps with the same agent and course, during which data was recorded. If necessary, after the practice lap, obstacles were replaced in the case they were pushed around, and the agent was reset to a "very safe" state. From pilots, I found each lap of my courses took around one to five minutes, depending on participant skill. I found participants took around two to three extra minutes on their practice laps as well, resulting in roughly 15 to 54 minutes of driving per person.

Before the two laps where data is recorded in each condition, I explained the distractor question to participants, so they knew to pay attention to the agent. After the laps were complete, the distractor question and other post-condition questionnaires were administered. The next obstacle course was prepared, the new on-screen interface (interactive agent or baseline) was explained, and the participant was again given a practice lap before continuing.

At the end of the experiment, participants were given the post-experiment questionnaire (interface ranking, final comments), and brought to see the course and robot in person. The details of the experiment were explained, as well as why I was purposefully vague on how the agents each conveyed the safety information. After any questions were answered, the experiment was over.

5.4.5 Implementation

My robot was a Clearpath Jackal robot running ROS Indigo. It was limited to 50% of its maximum forward and backward speed, and 75% of its maximum turning speed as pilot testing showed my robot moved too quickly in my smaller environment. A PointGrey Flea3 camera was mounted near the front of the robot such that the robot itself was not in the view of the camera. The camera was set to a 640x480 resolution (Figure 5.1b) at 45 frames per second over the institution's Wi-Fi network. The data handling and networking was handled through multi-threaded python code.

Participants were seated at a desk and allowed to adjust the setup to be comfortable. They used a 4K 27-inch monitor, with the interface maximized (black bars were used for letterboxing). They controlled the robot with a joystick (Microsoft Sidewinder USB) on the desk in front of them. The client-side was programmed in C#.

In my pilot studies I found that the robot was able to move my obstacles easily, making collisions appear to be not harmful to the robot and reducing the perceived negative consequence of hitting the obstacles. To make the obstacles more stable, they were each weighted with 14 KG of weights, placed on rubber friction mats, which in turn were placed on carpet tape stuck to the ground. With this much resistance, the robot could not easily push obstacles out of the way: operators needed to navigate the obstacle course correctly. Further, hitting an obstacle would usually stop the robot, produce an audible noise, and sometimes vibrate the cameras. To further emphasize collisions, my system would make the whole screen flash red briefly (1/3 of a second) when a collision was detected.

5.4.6 Analysis

I investigate the two components of my affective feedback strategy: a) how the agent behavior impacted operator mental state (if I induced an emotional response), and b) how this impacted the operator's driving behavior and teleoperation experience. The emotion analysis included the five emotion-susceptibility questionnaire subscales as covariates, to control for how people are affected by displays of emotion differently.

For performance measures, I analyzed collisions over time (number of collisions divided by completion time, in minutes), perceived workload (TLX sum and its subscales), my perceived safety and informative scales, and the safety rating calculated by my system.

Results

I recruited 23 participants by advertising with posters around my local university area. One participant did not complete the experiment, resulting in 22 participants (mean age of 28, standard deviation of 10.7 years; 10 female).

To understand how my interfaces may have changed operators emotionally, I ran a repeated-measures ANOVA on participant self-report measurements for valence and arousal changes, with the five emotion-susceptibility questionnaire subscales as covariates. I found a statistical effect of the interface on self-reported valence (a measure of pleasure or displeasure, $F_{2,32}=4.1$, $p<.03$, $\eta^2=.20$), and arousal (a measure of activation, or sleepiness, $F_{2,32}=3.4$, $p<.05$, $\eta^2=.18$). Post-hocs with Bonferroni correction found the daredevil agent produced higher self-reported valence than the anxious agent ($p<.02$, mean difference=-.32 points, 95% CI [-.59, -.05]). Other pairwise comparisons were non-significant.

Marginal means showed the daredevil case had the highest self-reported valence (mean=0.36, 95% CI [-0.07, 0.80]), followed by the numeric case (mean=0.27, 95% CI [-0.24, 0.79], with the anxious agent (mean=.05, 95% CI= [-0.44, 0.53]), having the lowest self-reported valence. For self-reported arousal, I found the numeric interface had the highest (mean=-0.68, 95% CI [-1.34, -0.02]), followed by the daredevil interface (mean=-0.82, 95% CI [-1.35, -0.29]), with

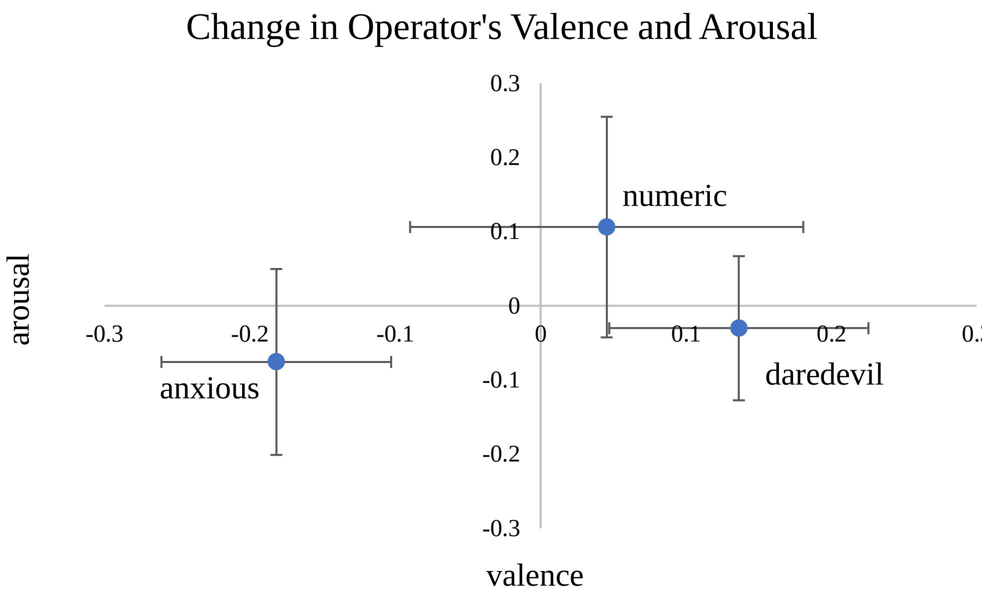

Figure 5.7. The reported emotions of operators (not their perception of the agent) after using each interface. Anxious interface appeared to lower valence more than the numeric case, while daredevil had lower arousal than the numeric (from contrasts). Grand mean differences are p<.05. Error bars show standard error.

the anxious case have the lowest self-reported arousal (mean=-0.86, 95% CI [-1.30, -0.42]). See Figure 5.7 (note for legibility, I have enlarged a sub-graph, but the scale ranged from -3 to +3).

There was an interaction effect between the subscale on susceptibility to happy emotions and the interface ($F_{2,32}$=4.4, p=.02, η^2=.22). I present the graph of the interaction (Figure 5.8) but note there were too few participants for my susceptibility scale, leading to few participants per valence rating, so I caution drawing conclusions from it.

I performed repeated measures ANOVAs on the performance measures listed above. The effect of the interface on collisions over time (CPM) was non-significant ($F_{2,42}$=2.6, p=.085, η^2=.11). Marginal means showed the numeric case had the most mean collisions (mean=1.9 CPM, 95% CI [1.6 CPM, 2.2 CPM]), followed by the anxious agent (mean=1.8 CPM, 95% CI [1.5 CPM, 2.2 CPM]), with the daredevil agent interface having the fewest (mean=1.6 CPM, 95% CI [1.3 CPM, 1.9 CPM]) – see Figure 5.11.

To test if the interface may have improved operation over time, as exposure to the reactions potentially affected driving behavior as time passed, I conducted a 2x2 ANOVA (interface versus time), with sample points at 10% intervals throughout the experiment. This was not found to be significant (p>.05) – see Figure 5.9.

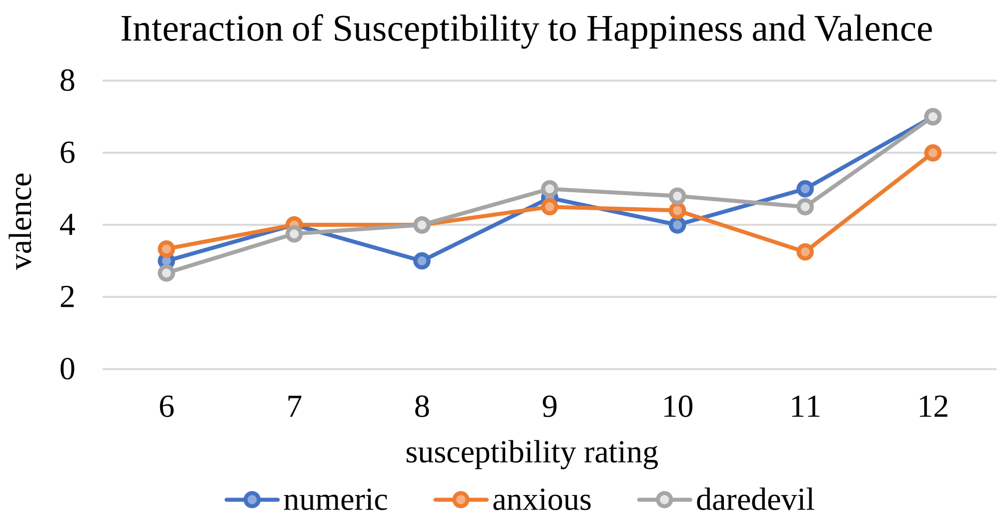

Figure 5.8. The interaction of an operator's susceptibility to displays of happiness, measured by questionnaire, and valence, by interface. The interaction is significant (p<.05), but we caution the number of data points per rating is small.

The agent's reactions may have been used by operators to inform themselves of their performance, but I found no statistical effect of the interface on perceived performance ($F_{2,42}$=2.7, *p*=.08, η^2=.11) – note the TLX performance scale is reverse-coded and higher scores mean worse perceived performance. The anxious agent interface made participants feel they performed worst (mean=9.4 points, 95% CI [7.5 points, 11.3 points]) – see Figure 5.10.

The perceived informativeness of the interface had a statistical difference ($F_{2,42}$=3.9, *p*=.03, η^2=.16) Marginal means showed the numeric case was perceived to be the most informative (mean=14.9 points, 95% CI [12.8 points, 16.9 points], followed by the anxious agent interface (mean=13.9 points, 95% CI [11.7 points, 16 points]). The daredevil agent interface was considered the least informative (mean=12.8 points, 95% CI [10.3 points, 15.3 points]) – see Figure 5.11.

All other tests and interactions were found to be non-significant.

5.5 Emotions, Performance, and Social Agents in Teleoperation

My results found differences in self-reported valence and emotion after using my interface, implying that my operators' emotions did change somewhat for each interface. I found

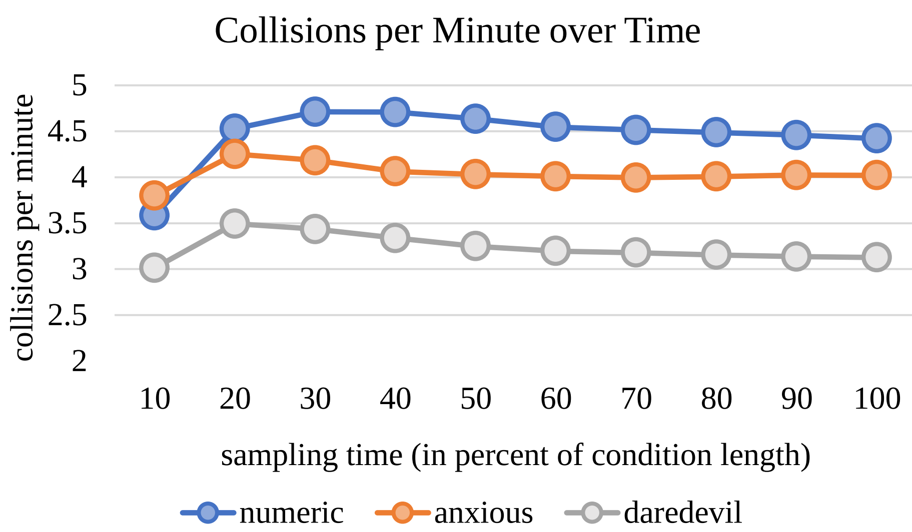

Figure 5.9. Collisions per minute sampled at 10% intervals through the study. High variance in our sample means we could not detect a difference (p>.05).

inconclusive evidence for the numeric case to have a higher collision per minute rating than the social interfaces, and that collision rates may be stable over time, though future study is required to confirm this. Although I did not detect a statistical difference, my results indicating potential differences merit further inquiry. Daredevil had the lowest average collisions per minute and was perceived by operators as enabling the better performance of my three interfaces. The anxious interface had fewer mean collisions per minute than the numeric case, but was seen as having worse performance. The numeric interface was the most informative, followed by the anxious and then daredevil interfaces.

The changes in valence and arousal demonstrate that an on-screen agent using affective feedback of safety ratings can change an operator's mental state. When inspecting average safety scores, I found that people, on average, drove in a way that my system rated as unsafe. Thus, people would have seen primarily a negative reaction from the anxious agent, and a happier face for the daredevil agent. This aligns with my background theory and results: the anxious interface (a lower valence emotion than happiness Posner et al. 2005) was reported as making participants feel lower valence overall, and happiness (a higher valence emotion Posner et al. 2005) had a higher valence. Thus, I can see the expected emotion divide (happy, sad) between daredevil and anxious, acting as a manipulation check that viewing emotions in a teleoperation interface can influence the operator's emotions to become more similar to like the displayed affective behavior. While the differences were small, I note that the interaction overall was very short. Small differences over time, however, may amount to a longer-term effect, but more research is needed to confirm this.

Our theory that positive emotions would encourage the behavior at the time of the affective feedback was not supported by my data. It is possible the emotional response itself was not strong enough for this effect to take place. Another possibility is that the daredevil persona helped people relax; when colliding, the reaction on the face was happiness, which may have reassured the participant. If they saw the anxious persona look unhappy and experienced my observed negative valence shift, instead of discouraging the behavior, the feedback may have made them tense up and perform worse. This may also explain why self-reported performance was higher for the daredevil persona: the positive reaction upon mistakes made participants think they were not doing poorly. Certainly, the intricacies of how participants reacted to the emotion needs further research for clarification. Further, this highlights the importance of a more rigorous model for creating personas, which would enable me to more concretely and specifically reflect on components of the agent's reaction.

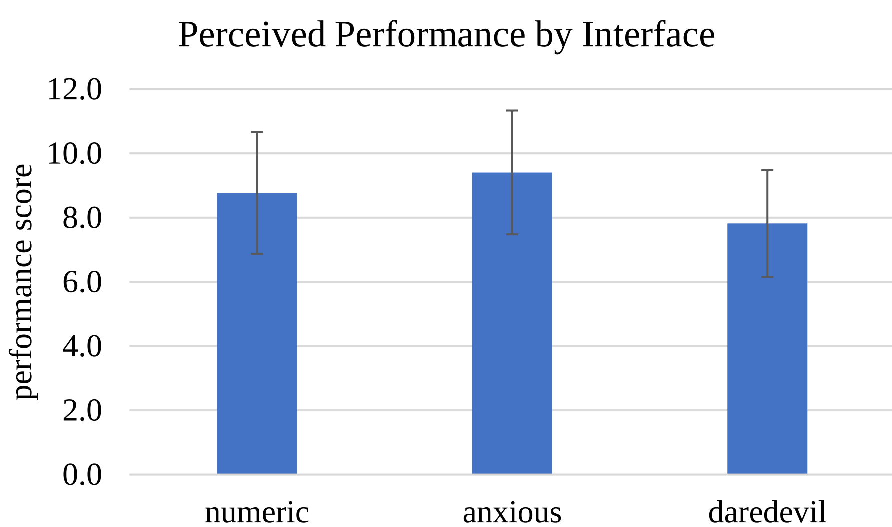

Figure 5.10. Self-reflection performance values by operators were not significant (p=.08). Performance is reverse-coded (higher means worse perceived performance).

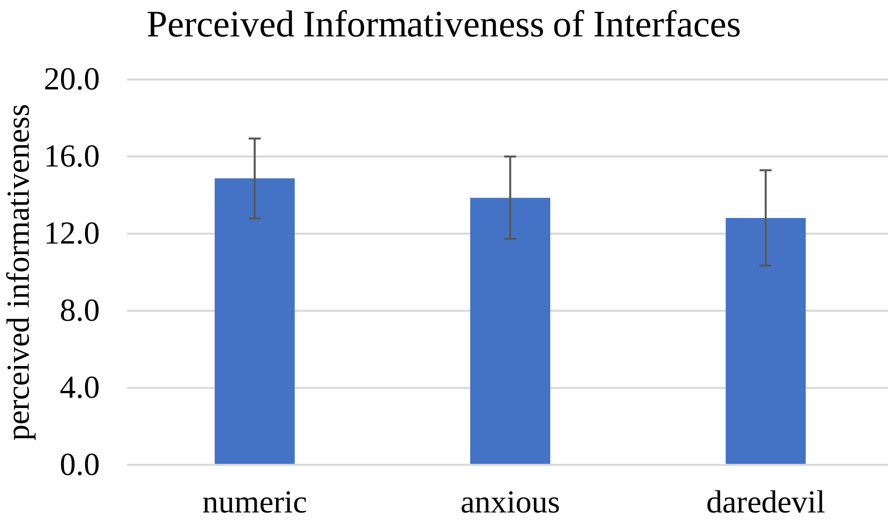

Figure 5.11. The perceived informativeness of the interfaces was different (p<.5). Error bars show 95% CI. Interestingly, the difference with the numeric case was slight, despite very different visualizations.

Even though my collisions over time were not found to be significant, I wish to briefly discuss the interpretation of my descriptive statistics. I observed a non-significant difference in collisions over time of 0.3 collisions between the numeric and daredevil case. Even over a short period of operation, such as 30 minutes, a difference of 0.3 collisions per minute results in an extra 10 collisions. Further, as per Figure 5.9, the average difference in collisions per minute between interfaces may be stable over time. While I cannot state that my social interfaces affected driving safety, even if small differences can be found I believe it would be worth further research.

As a related aside, the numeric interface appeared to have higher collisions per minute. This may be due to the mental work needed to read and understand numeric data, which may take more attention away from actual operation. However, I stress again that these collision results are not statistically significant, and further study is needed to confirm if my measures reflect the population correctly.

The daredevil interface was seen as least informative, which may be due to it being unintuitive: after seeing a more positive face after a collision, operators may have thought the system was not working properly. However, all three interfaces were ranked similarly (Figure 5.11.), which may suggest that social interfaces for communicating information may be feasible when the operator does not need a granular understanding of data, such as in my case.

While I motivated my design with existing social psychology theory, social interfaces have sometimes struggled in industry. Microsoft's Clippy is one example: critiques of Clippy point out that Clippy breaks social rules, such as offering help when it is not asked for and not remembering people's decisions (Whitworth 2005). Others have argued virtual companions should be more agreeable (such as my daredevil showing happy faces during bad driving), or offer alternative solutions to a difficult task (Luria 2018; Vázquez et al. 2014). My agents were only reactive and did not try to provide advice to users; this may be why I did not witness similar negative feedback to my agents. It is possible that virtual agents helping in teleoperation may have different social rules applied to them and is an important avenue for future work.

5.6 Complexities of Evaluating Social Teleoperation Interfaces

My results raised many questions for future work. Overall, I aimed for external validity and used an algorithm that could be applied to robots right now, but that made it difficult to evaluate the nuances of the social interfaces. For example, a future study could look at if the agent even needs to properly react to current driving behavior: it could always look annoyed, or happy. This would remove the variable in my studies where drivers of different skill levels would see, on average, different agent reactions. If certain agent reactions would have a stronger effect than others, I would have difficulty measuring those effects as my operators had different levels of exposure to each reaction.

The agents were described primarily as a tool, or algorithm. It is possible that this reduced the anthropomorphism effect and reduced the impact of the agent's reactions. If operators thought the agent was intelligent, it is possible they would react to the agent in a more social way. Further, the agent could be presented in multiple different ways: as a boss, as a coworker, as the robot's intelligence, etc. This change in agency and relationship with the operator could further affect their reactions to the agent's displayed emotion.

The operator's initial mindset likely also affects how they accept the agent as a social entity. For example, games can leverage suspension of disbelief – players are willing to believe a character is real because they are there to be entertained. Teleoperation may not currently have this benefit as it is real-world task oriented. However, much like my inspiration was taken from video games, other narrative techniques could be used to introduce or continually interact with the agent could be employed to further the emotional connection with the operator.

The emotion susceptibility questionnaire I used had five subscales, which proved difficult to use with a smaller participant pool. With my participant numbers, adding all scales as covariates may lead to overfitting for my model. Due to the direct link between emotion susceptibility and viewer response in the literature, I believed including the covariates for my emotion data (valence and arousal) was necessary. However, I opted to use a simpler model in my performance analysis for fear of overfitting.

I witnessed a large amount of variance between participants. This may be due to my course being overly difficult, such as having little room for the robot to make turns and operators having difficulty visualizing the robot in the remote area. Anecdotally, I witnessed operators who performed well, but got stuck in difficult situations on occasion, resulting in numerous collisions during a single event, perhaps increasing the variance in my results. While I pilot tested extensively to calibrate my agents, I still believe that the course difficulty may have confounded my results. I recommend future work carefully consider and calibrate the difficulty of their study, and reconsider their measures of unsafe driving.

My choice of baseline may contribute to my results. I opted to design my baseline to have information parity with my affective feedback agents: all interfaces, on some level, presented collision and velocity information. However, by displaying safety information numerically, I possibly increased the mental processing needed to understand the presented information as compared to the affective feedback from my agents. To reduce this, I could have a baseline with no information displayed, or just a neutral face displayed, or to use simple text labels for the emotional state (e.g., "safe," "unsafe," "very unsafe," etc.). Thus, my baseline is not a truly neutral control condition, but enables me to compare social to non-social interfaces without the confound of difference in available information.

One limitation I believe is the presence of the researcher in the room while the participant was piloting the robot. As the researcher was an authority figure and a stranger, the operator may have suppressed their reactions to appear more professional and under control to the researcher. The researcher's own subtle and subconscious body language may have provided a stronger social signal to the operator than the agent itself. Thus, I recommend removing the researcher from the room in future studies.

5.7 Towards Social Interfaces for Supporting Teleoperators

I presented a proof-of-concept for using affective agents to shape teleoperation experience and highlights the potential for using social video game and social HRI techniques between the teleoperated robot and the operator. My results highlight using social interfaces similar to some found in video games can leverage affective feedback can indeed change a teleoperator's mental state, and its design impacts how effective they felt the feedback was for communicating state. Specifically, I found I could change an operator's emotional state, one important part of interaction design (Chapter 2). Emotions are an important part of the user experience, and may itself be correlated with driving safety in motor vehicles, and so I recommend even further and more targeted research into affecting an operator's emotional state through design.

Unfortunately, I did not observe any other changes in teleoperator experience or behaviour. I originally expected that there may be some small effect, as suggested by the related work. However, the small successes I found have helped me demonstrate the potential of, to my knowledge, a new avenue for teleoperation interface design research, melding knowledge from teleoperation, social human-robot interaction, and video games.

The technique developed in this chapter was inspired by social displays characters in video games. While learning from video games to improve robot control and design is important, my results demonstrate how video games may be able to influence operators in other, broader ways. Video games, as an entertainment medium, often work on social, mental, and generally playful levels that are not purely focused on function or performance. Using such a strategy for emotional effects, I changed how the operator feels after using the robot, and focusing on those other game design goals may further be able to affect teleoperators on an emotional or otherwise experiential level.

CHAPTER SIX:

A FRAMEWORK OF VIDEO-GAME INTERACTION TECHNIQUES FOR TELEOPERATION APPLICATIONS

Through the years, the video game industry has iteratively designed and developed a number of interaction paradigms to improve usability and user engagement, and in this thesis I have proposed that teleoperation research can leverage this knowledge to improve operator performance and experience. In earlier chapters, I explored a number of video game-inspired interaction techniques and experimentally verified their impact on teleoperation performance and experience. In light of these successes, I aim to more broadly and generally learn from video game interaction design for improving teleoperation.

In this chapter, I present an original framework of video game interaction techniques that classifies and abstracts video game interaction designs from a number of angles to provide the tools and vocabulary to discuss the application of broad video game techniques to teleoperation. I detail my data-driven approach and analysis methods I used to create this framework, and ground my results at each step in how those results and my framework apply to teleoperation problems. I further draw links between my framework and existing solutions in the broader literature, and use my framework to discuss potential new solutions to teleoperation interaction design problems.

My analysis focuses on organizing video game interactions by identifying the general goals and strategies they use to shape player experience. To understand these general approaches, I survey a number of critically acclaimed games and observe what interaction designs they employ and what problems those designs aim to solve. By focusing on the higher-level goals and design strategies in games, I am able to build a framework that classifies interactions into abstract groups that can be used to draw general parallels between video game and teleoperation design based on their problem and solution spaces.

Throughout my process, I ground each concept in my framework back to known problems, needs, and solutions in teleoperation interaction design. This provides broader evidence for the applicability of my approach outside of the examples I myself developed. It also helps me understand what areas of video game techniques are still unexplored in teleoperation, providing directions for future work. Overall, this chapter provides a more general and theoretical grounding to my approach, describing the similarities between video game and teleoperation interaction design problems and solutions, and highlights the potential for leveraging further knowledge from the video game industry in future telerobotics designs.

6.1 Why Build a Framework of Video Game Interaction?

Throughout this thesis I have argued that teleoperation and modern video games share many interaction design problems, and that teleoperation may be able to take inspiration from video game interfaces to improve teleoperation user experience. Using the projects in this thesis as a base, I ask if my approach of learning from video games can be applied more broadly to improve teleoperation interaction design. However, there are many games to investigate, each with potentially interesting or unique interaction designs that may be useful to teleoperation. Instead of considering each of the thousands of games and interactions individually for its applicability to teleoperation, I aimed to distill this variety into a more compact and simple form that can be used to explore and reason about video game interaction broadly in its applications to teleoperation.

One common solution for understanding or simplifying a complex problem space is to develop frameworks or taxonomies. These are useful tools that take a large design space and broadly group different parts of a problem, classify common solution approaches, create vocabulary for describing and discussing phenomena, or provide evaluation tools. For example, in human-computer interaction, researchers have outlined different design dimensions that can be leveraged when creating and evaluating learnable (Grossman, Fitzmaurice, and Attar 2009) or engaging (Langer, Hancock, and Scott 2015) systems. In this latter work, the authors outline three different dimensions of suspense and use this vocabulary to describe which situations should use certain types of suspense; specifically, they use their framework to analyze and classify traditional software tutorial designs using the authors' identified dimensions, which enables the authors to identify new approaches to design tutorials by exploring their dimensions in ways not covered by traditional approaches. Other examples of taxonomies and frameworks have been used to discuss how others can evaluate and classify user sentiment about a system (Nasukawa and Yi 2003), and have classified existing work to identify successful approaches and outline new research directions in broad interaction domains such as tangible interfaces (Ullmer and Ishii 2000) or computer input devices (Buxton 1988). I follow this practice with my aim of providing an overarching framework to help describe the design space of video game interaction, highlight the coupling between video game and teleoperation interaction, and provide researchers with a toolkit that they can use to explore video game designs to help with a given problem in teleoperation.

6.2 Approach and Contribution

To create my framework I take a qualitative approach drawing heavily from grounded theory (Glaser and Strauss 2017), making observations of data, applying descriptive codes to the data, grouping and relating the high level codes, and iteratively refining and writing the framework. To gather a broader variety of interaction techniques in games to analyze for my framework, I perform an exploratory survey of critically praised video games and make observations of their interaction techniques. I consider the large design space of video game interaction techniques from a series of angles. I methodologically break interactions down into their components (e.g., commands from a user, feedback from a computer), and through qualitative analysis methods aim to identify underlying archetypes of interaction techniques and methods used in games.

Through all of this, I focus on what purposes the video game techniques serve, general strategies used in implementation, and design goals, while continuously grounding my results in their applicability for teleoperation.

Resulting from this process, I created an original framework of video-game interaction techniques grounded in the needs and problems of teleoperation. With it, researchers can use the framework as a language for describing, comparing, and contrasting video game interfaces for teleoperation at a higher level, and can use as inspiration for developing additional novel teleoperation techniques. Specifically, I discussed the core information that video games and players exchange, how they exchange it, video game design goals, general implementation strategies, and specific interface techniques. I examined how each of these angles relates to common teleoperation problems and solutions. Linking these components together, my framework provides one approach for understanding and discussing video game interaction design and their broader use in future teleoperation designs.

Throughout this process and from my resulting framework, I provide the human-robot interaction community with a bridge from the video game industry to teleoperation. While the works earlier in this thesis act as deep, experimentally supported designs of video game inspired teleoperation interfaces, this research acts as a theoretical grounding for the broad applicability my approach. I demonstrate that games have solved numerous other problems related to teleoperation, and that some techniques, outside of this thesis, are also seeing success. Learning from video game interaction techniques is a viable approach for improving teleoperation.

6.3 Methodology

My approach for constructing my framework comes in three parts. I survey video game interaction to gain a broader sense of the techniques and approaches video games use. I then perform qualitative coding on my observations to organize them into high level groups and themes, reducing my observations into a simpler and more abstract form. I analyze these groups with an interaction framework from human-computer interaction to provide deeper insight into how video games approach interaction, and further iterate and use qualitative coding with these new results, structuring these new higher level groups to produce the framework itself.

6.3.1 Survey Methodology

To collect data, I played and watched my selected video games while thoroughly documenting the interface techniques, and messages used. I outline my methods and motivations for selecting games and how I made my observations below.

Making Observations of Interactions in Video Games

I wished to analyze a number of video games, but thoroughly playing a video game is a resource-intensive task. Thus, to collect data, I both played games as well as watched recordings of other people play games via online video services. Watching and playing each make certain observations easier; while playing, it is easier to understand how controls interact with the game: an observer may not know what buttons are being pressed while watching videos of gameplay. In contrast, observing games frees the watcher from any exertion and concentration required to play the game, and can focus their attention freely on the interface. Further, observers watching a video can freely move ahead in time, skipping repetitive gameplay scenarios and long tutorials to view a wide variety of gameplay that may otherwise take days of play time to see. I set a goal of playing around 50% of the games in my dataset and observing the rest; this strikes a balance in having the researchers gain experience controlling many different types of games, and freely observing without any worries about gameplay or game performance.

Selecting Video Games Similar to Teleoperation

There are thousands of games that can have multiple different interfaces used in different situations over dozens of hours of gameplay each. Thus, my sampling approach was to filter existing games with rules that would select games similar to teleoperation; not all games are applicable to teleoperation – one may be hard pressed to improve teleoperation with interfaces from a puzzle game, such as Tetris (Pajitnov, A.L., 1984); I narrowed my search to games that share similarities with teleoperation: typically completing objectives in a 3D world by controlling one or more avatars. Even with this filter, I can include games from a variety of genres including adventure games, role-playing games, first-person shooters, real-time strategy games and more.

I do not suggest that video games have universally good designs; games, just like any other software, have numerous bad examples of interface design. To mitigate including such interfaces in my survey, I decided to focus on games rated highly on review aggregate websites (e.g., Metacritic, www.metacritic.com/game). Metacritic is an independent website that aggregates game review scores from multiple publications and averages them, giving an overall sense of critical reception. And while there are many highly rated older games, I further assume that video games have been improving interface designs over time, and thus I focus on more modern games from the aforementioned lists. While not completely representative of the sample space of games, I note that the sample space of games is very large, and I aim simply to provide a base sample for my analysis to extract insight into the applications of game techniques to teleoperation, which can be extended afterwards with the addition of more games.

Of the games selected from review websites, the games actually played (as opposed to observed) were a convenience sample based on price and availability. This is not as desirable as a random subset of the well-reviewed games being played: it may produce sample bias and will limit the generalizability of my results. However, due to my exploratory goal of observing a variety of interaction techniques in video games, the results of the types of interfaces I see and how they work will still be relevant: I am not making arguments about which types of interfaces are used more or less, for example.

6.3.2 Analysis Method to Create our Framework of Video Game Interaction

My game survey resulted in rich and complex observations, so I employed qualitative analysis techniques to simplify and relate potentially disparate interfaces into groups of interfaces that share goals or techniques. With these techniques, my end goal was to relate those groups of interfaces in a meaningful structure that enables designers and researchers to understand and reason about the design space of video game interfaces on a more abstract level than on a per-interface basis. My methods were based on grounded theory (Glaser and Strauss 2017) with open coding (Berg 1989; Moody 2009; Andersen 2001).

My approach involves making observations of complex phenomena to generate data, creating codes to describe the data, and iteratively refining and rewriting results until it describes the dataset (Glaser and Strauss 2017). Open coding is one way to generate *codes* for complex data – words or phrases that describe the theme of a group of observations (Corbin and Strauss 1990). For example, I may create codes for a health bar in multiple ways, such as "displaying health on-screen," and "representing health as bars." This process is iterated on, updating codes, combining them, and removing them from the code pool (Martin 1986; Corbin and Strauss 1990). With the simple example codes above, I may compare observations with other similar codes in the data, such as seeing what other types of information are displayed on-screen, or displayed as a bar. Open coding does not rely on developing codes before analysis, but lets them emerge from observations: researchers tag data with conceptual labels they create, and constantly compare those codes to others used in the dataset (Corbin and Strauss 1990). This analysis of the codes themselves may lead to the creation of new codes to explain the data along new dimensions, a process known as axial coding (A. L. Strauss 1987; Corbin and Strauss 1990).

Axial coding is useful for the goal of identifying emergent relations between codes, uncovering new codes, or identifying categories or subcategories of codes. I combine axial coding with open coding to discover nuances within themes, creating more detailed grouping of the data (A. Strauss and Corbin 1990; A. L. Strauss 1987), and is one way to create "sub-concepts" from the original conceptual code (Corbin and Strauss 1990). In my earlier example, I may see many codes that deal with representing health or other resources (such as remaining energy); here I may axially code a parent code "displaying depletable resources" with sub-concepts that deal with graphical representations (e.g. a bar), or precise numerical descriptions (e.g. stating 56% health remaining). Thus, open coding and axial coding help my goals of abstracting groups of interfaces into high-level common interaction designs.

To create a framework from simple codes, I iteratively abstract my codes to create high level categories which can be linked by how they interact with the other categories. Low-level codes are combined into concepts that help explain high-level similarities between many codes (Martin 1986; Corbin and Strauss 1990), such as a concept for displaying player status instead of a codes for displaying each type of player status (health, energy, etc.). Similar concepts may then be combined into categories (e.g., communicating properties of an in-game object). By describing the links between these categories, I create a framework of video game interaction grounded in my observations. This iterative approach helps me create the framework by linking all my observations through multiple layers of abstraction. At every stage of abstraction in my results (the codes, concepts, and categories), I identify archetypical interaction designs and strategies to apply to teleoperation. These abstract classes provide the vocabulary to reason about and discuss groups of interfaces, and can provide inspiration for teleoperation design by looking at specific instances of those interaction classes in my framework.

6.3.3 Using A Message-Passing Paradigm to Analyze Interactions

There are many dimensions to consider when analyzing interaction techniques in video games. For example, an interaction technique may be communicating information to a player (e.g., player health), can have aesthetic design elements (e.g., to make it look archaic or futuristic), and can use a mix of modalities (e.g., visual, aural, etc.). Further, there is a constant exchange of interactions between the player and the game – the player will be controlling a character in-game while the game simultaneously provides many types of feedback (Figure). To help me simplify this complex web of interaction and communication into simpler and more manageable interface and interaction design components I employ a message passing paradigm.

Figure 6.1. In Destiny 2 (Bungie, 2017), a player is sending command messages simultaneously with the mouse and keyboard to control both movement, character actions, use of weapons, and UI manipulation. At the same time, the game is sending many feedback messages about objectives, player state, navigation, items and rewards, and more. Our message passing framework enables us to analyze each element on a simpler scale, while also enabling us to group similar interface

The message-passing paradigm is an analytical paradigm used in human-computer interaction as one way to understand interaction by providing a framework for helping identify and focus on particular components of interaction (Andersen 2001; Moody 2009; de Souza 1993). In particular, I use this approach to separate what information is being conveyed (e.g., a character has low health) in an interaction from how that information is presented (the interface or implementation, such as a loud aural alarm). This separation enables me to compare specific techniques at an abstract level, understand which interactions convey the same information, different strategies for conveying the same information, and how the different implementations are designed to affect the user experience.

Message Passing to Analyze Video Game Interaction

Message-passing is a paradigm that views interaction as a person and a system (e.g., video game or teleoperated robot) communicating exclusively by sending messages back and forth: the person gives commands (messages) to the system, and the system communicates status updates and results (messages) back to the user (de Souza 1993). This can describe both asynchronous interaction (e.g., the user enters a command, and the system communicates the results to the waiting user), or synchronous interaction (e.g., the user manipulates a joystick to control avatar speed, and the system constantly provides visual updates to the user). Because of my specific application of video games and teleoperation, I define a message sent from the user to the system to be a *command message*, and when the system sends a message to the user I call it a *feedback message* (Figure 6.2).

In mesage passing, *messages* are the core abstract information that one is trying to communicate. However, for a message to be conveyed it must go through two important steps. First is *encoding*, where the message is embedded within a representation – an interface that

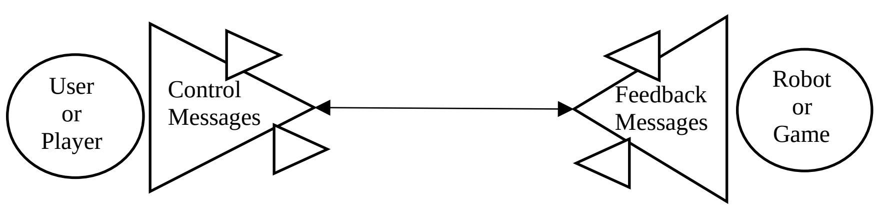

Figure 6.2. A high-level view message-passing in our video game or teleoperation scenario: a user and a system sending and receiving messages constantly in real time.

defines how the message can be interacted with. Second, is *decoding*, where the encoded message must be interpreted. The encoding and decoding are how the user or system engages with the representation of the interaction (de Souza 1993). For example, if a user wants to direct an avatar to move forward (*message*: avatar, move forward) it must be encoded (*encoding:* user pushes joystick forward to indicate a move command) and subsequently decoded (*decoding:* system translates joystick offset to an avatar action). Following, if the system wants to convey the new location to the user (*message*: avatar's updated location) it decides how to convey that information (*encoding:* move graphical representation on screen), which the user must interpret (*decoding:* user must recognize and understand visual change in avatar position on-screen). These components are visualized in Figure 6.3.

I apply my message passing paradigm to analyze interfaces identified in my qualitative analysis. Within each code or group of codes, I extract the core information (the message) and encodings (the implementation) of the interfaces and then perform iterative coding on messages and encodings. Thus, I integrate message passing into my analysis to give me a deeper understanding of interfaces during my open coding process.

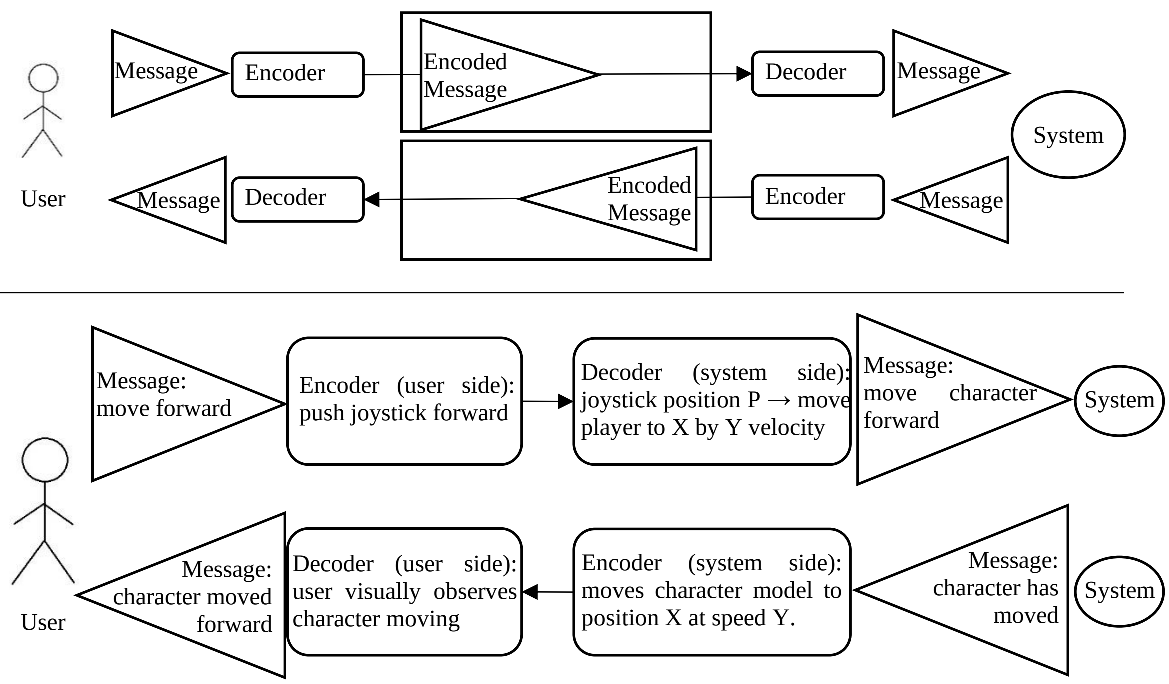

Figure 6.3. (Top) The message passing paradigm for HCI, adapted from (de Souza 1993).

(Bottom) Here, a high-level message is passed from the player to the system, described roughly with a message passing paradigm adapted from (de Souza 1993). *Left-to-right, a user commands their character to move forward by pushing a joystick (a command message). The bottom row shows the system showing the result of the command (a feedback message). In each case, the core message must be encoded and decoded in order to communicate the message.*

At first glance, this may seem to be a cumbersome approach that adds complexity to a simple idea (interaction). However, message passing provides more abstraction and analysis power than when I simply consider specific instances of interaction techniques. In my case, this enables me to move beyond considering particular game techniques (e.g., a specific "player health" visualization for use in teleoperation) and toward more abstract, generalized inquiry (e.g., games use a range of visual encodings to communicate player status). Message passing enables me to separate communication intent (messages) from the interface (encoding) and interpretation (decoding), and to deal with each individually.

The concepts of messages, encodings, and decodings further help me reason about applying a video game technique to teleoperation. By looking at the message independently of its encoding (e.g., player status rather than just considering a health bar), I can understand which teleoperation problems it may apply to (e.g., displaying robot status). Studying encodings helps us understand different designs for presenting a message to users. Decoding then helps us understand what the design challenges are – that is, understand how a user may have difficulty interpreting an encoding. The separation between the underlying information being communicated and how it is conveyed enables us to consider useful encoding strategies (e.g., visual techniques, input devices) for feasibility and usability, independent of what message is being conveyed.

This message passing concept – messages that are encoded on one side and interpreted by the other party through a decoding process – forms the backbone of my analysis into leveraging interface techniques used in video games.

6.4 A Framework of Video Game Interaction for Improving Teleoperation

In total, I surveyed 30 games from a variety of platforms (personal computer games, various television game consoles, see Appendix D). One researcher played or observed all games for 30 to 60 minutes each (playing 26 of the 30 games observed), with the researcher making, on average 28.5 observations per game for a total of 853 observations.

The primary result from my survey and qualitative analysis is a novel framework for describing and understanding the various components of interaction in a video game. This framework provides two important contributions: one, it separates interaction into a small number of core components: the communication purpose, the communication method, higher level implementation strategies, and broad user experience goals. On another level, I discovered that the vast range of video game techniques could actually be described using a small set of archetypes within each core component of my framework, such as how only a small number of messages can describe most interactions I observed. While I detail each framework piece in great detail below, here I highlight each high-level core component and provide examples in Figure 6.4. Each component and its underlying archetypes provides a new perspective on how to view interaction between the user and the system.

Archetypical Messages – My analysis revealed that the majority of information communicated between the player and system can be described with a relatively small set of general messages. For example, an environment awareness message may alert a player to an important item behind them.

Encoding Parameters – While the specific ways that games encoded messages varied immensely, underlying these a core set common design parameters shared across interfaces emerged from my analysis. These *encoding parameters* included elements such as when an

Messages	Encoding Parameters	User Experience Goals	Design Strategies
Information communicated between user and system	Design decisions for conveying a message	Overarching design goals direct design decisions	direct how encodings work together and support experience goals
e.g., player status moving a character directions to next objective	e.g., granularity of information when an interaction begins on-screen location of encoding	e.g., challenge aesthetic satisfaction	e.g., group similar encodings control simplification interface saliency

Figure 6.4. The top-level of our framework, with a brief description of each component followed by examples we found in our video game observations.

encoded message is sent to the player, how specific the information in that message is, or where the message is displayed (a position on-screen, a position in the audio landscape, or even in a controller).

User Experience Goals – My results included common design goals across a wide range of messages, encodings, games, and genres. These goals focus on different aspects of the game experience, and they acted as guiding principles for how and why a message was sent or a certain design was chosen. For example, a message and specific encoding may be designed to support high-level interaction goals, such as fun, challenge, or to be easy to learn.

Design Strategies – My analysis also found strategies common across many games that help direct encoded message design to work together in an interface to support a design goal. Generally, these strategies assist in helping numerous encoded messages work together in a cohesive interface. For example, the strategy of simplification can be used to enable a simpler interface, such as by removing unnecessary controls when appropriate (e.g., removing battle commands in a friendly city).

My framework provides multiple perspectives to understand video game interactions; for example, health indicators are a common interface in video games. These send a *property message* – the information that a player's health is at a certain value. This message may be encoded in multiple ways. If I want the player to know their health with precision (*encoding parameter: information granularity*), it may be displayed on screen with a number. If less precision is needed, I could instead represent a player having less health by vibrating a haptic controller (encoding parameter: *modality*), or only encode the message when the player's health changes (*encoding parameter: temporality*). This latter example helps implement a *temporal* design strategy which in turns supports the user experience goals of *understanding* (by conveying player state when needed) and *simplification* (by, perhaps, reducing visual clutter). With a different goal (perhaps *immersion*), a health bar may instead be represented with the virtual avatar looking physically more injured (encoding parameter: *diageticness*, design strategy: interface saliency).

Thus, each framework component is connected to other levels of the framework. Archetypical messages are encoded in message encodings that themselves are made up of encoding components; multiple encoded messages make up an encoded message group, which implement a design strategy. These strategies help fulfill a high-level design goal. In this way, my framework provides a high-level structure to view and understand interaction design in video games. I consider each level of the framework, as well as the archetypes of each level, in the context of its use to teleoperation. In the same way I can reason about and understand video game interactions from different perspectives, my framework further helps consider how different video game interfaces may apply to teleoperation by enabling the comparison of larger groups and archetypes of interactions.

I provide an overview of each framework component to provide a better sense of how each piece works together. Then, in each following section, I describe my framework components in detail.

6.4.1 Overview: *Archetypical Messages*

The concept of messages helped me identify a variety of common types of information passed in either interaction direction. I found I could group these messages into two large groups: messages to the user (feedback messages) including representations of the game world (the character(s) the user controls, the environment around those characters), conveying character status such as health or available items, or objective and navigation information. Messages from the user (control messages) inform the system of what the player wants to do, including character actions (movement, attacking, using an item), changing the system's state (opening menus, pausing, starting a battle), or changing the camera view to better understand the area around their character. Note these are all abstract and demand no particular implementation – movement could be dictated with voice commands, a joystick being pushed, or with a camera that reads user gestures. I found a small number of common messages could describe most of my observations – messaging passing helped me distill video game interaction to a number of basic messages..

6.4.2 Overview: *Encoding Parameters*

Many games share a base set of messages, but the implementation of these messages can vary wildly, leading to different gameplay and user experience. However, my analysis found a small set of design parameters common across all implementations I observed, and these parameters helped explain the differences between encodings. For example, a given control message can have different encodings: a move command encoded with dictated voice commands [Mass Effect 3, BioWare, 2012; Hey You, Pikachu!, Nintendo, 2000], a joystick being pushed (see most console games), or with a camera that reads user gestures [Kinect Star Wars, LucasArts, 2012]. These example interfaces send the same command, but vary in their *medium* (aural, haptic, or visual respectively), one of my encoding parameters. Thus, encoding parameters serve as vocabulary that highlight elements of an interface implementation that can be explored if a new or better design is being developed. This vocabulary further acts as a high-level classification framework that can help discuss and reason about entire groups of interface implementations that share similar encoding parameters.

6.4.3 Overview: *Design Strategies*

Our analysis also found strategies common across many games that help direct encoded message design to work together in an interface to support a design goal. Generally, these strategies assist in helping different encoded messages work together in a cohesive interface by acting as heuristics and guidelines for encoding design. For example, one strategy could be to reduce the visual saliency of interface components when other information becomes more relevant to the player, or to combine encodings for similar messages into compound interface elements, such as grouping all property message encodings for a character in one interface element. These overarching strategies may be applied to a variety of techniques: healthbars, warning messages, mission objectives, navigation aids, and more. I explicitly consider how these strategies could support design goals in video games, and how they may mitigate or solve

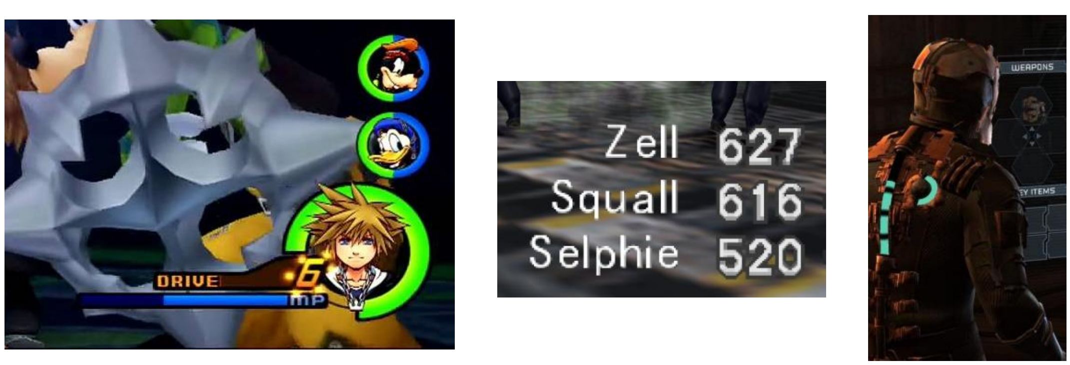

Figure 6.5. Comparing the health indicators of Kingdom Hearts (left), Final Fantasy 8 (middle), and Dead Space (right). Each demonstrates a different encoding of the same message: remaining health. The encoding can alter player experience because of the granular presentation, how embedded it is in the game world, when it is displayed, and more.

problems similar problems to those in teleoperation.

6.4.4 Overview: *User Experience Goals*

In addition to design strategies, I saw another category of codes arise from my analysis: user experience goals. These are high-level guiding principles that can direct all interaction design in a game towards creating a specific user experience. Once picked, goals can be used to guide and evaluate a design. For example, a new interface could be used to control a character's movement. If the design goals were to be have simple controls that create satisfying in game animations or results, a designer can reflect on the interfaces and ask, "Do these interfaces support my user experience goal?" One could further imagine a designer wanting to increase player uncertainty to promote cautious play. The designer may choose not to send the message of maximum amount of health (e.g. Figure 6.5, middle), and only show a shrinking health bar without an outline. Thus, the interaction design (in this case, *not* sending a maximum health message), supports the goal of making the player uncertain. This makes it difficult for the player to understand how much danger they may be in, potentially changing the way the experience and play the game. I outline several goals I found repeated in my data, and consider how these goals can be used to guide teleoperation design as well.

6.4.5 Overview: Framework Details

In the rest of this chapter, I detail the messages, encodings, encoded message groups, design strategies, and design goals identified from my analysis of my game observations. At each stage I consider how the results can be applicable to teleoperation. I end the chapter with a high-level presentation of the framework and reflection on its use in teleoperation as a whole.

6.5 Archetypical Messages in Video Games

I begin my results presentation with the emergent message types I found throughout the video games I observed. These *archetypical messages* describe the interactions within the game, but do not specify how they are presented to my created by the player. In other words, these messages present a classification of common video game interaction ideas.

Despite a range of game genres with seemingly radically different interaction paradigms, my analysis revealed that the messages flowing between users and games fall into a small set.

Some message types had emergent, more specific submessage types – this breakdown provides further nuance and depth to the types of a single message type. These archetypical messages represent a large number of possible messages used in video games, and by seeing if a teleoperation interaction sends similar messages enables me to reflect generally on which video game interfaces may be relevant for teleoperation.

I discovered an asymmetry between the command and feedback messages – those sent to a system or to the user from the system – and so I present them independently below.

6.5.1 Archetypical Feedback Messages

Feedback messages transmit information from the system to the player, and I found seven high-level archetypical messages in total. I detail each further below, but present a high-level summary here:

> *Property* messages: the game communicates properties or state information about the player character and other in-game entities, such as well-being, abilities, resources, available actions, etc.
>
> *Interactable* messages: the player will need to interact with the environment; interaction messages communicate what the player can interact with (e.g. there is a door that may open), how to interact with those things (the door can be opened with a button press), and the results of the action (did the door open, or was it locked?).
>
> *Navigation* messages: the player needs to be sent navigation messages that tell them where they have been and where they need to go, e.g., "to reach the dungeon, head in

Figure 6.6. Three different messages all use the same control encoding in Dragon Age: Inquisition (BioWare, 2014). In each case, the user places the mouse cursor over an area and right-clicks. In the left case, the message is to pick up the item. The middle image sends the message to attack the enemy, and the right image shows the message being sent to move the character to that position. Since the decoding of the message encoded in the right-click is dependent on the game state (what is under the cursor), the game gives feedback messages (a different appearance to the cursor) to show the user how the system will decode the right-click.

this direction."

Environment Awareness messages: A player is sent environment awareness messages to be aware of things around them, such as dangers, other computer characters, treasures, etc. E.g., "there is an enemy nearby to the player's right." This differs from "Interactable" as it emphasizes what is happening around the player, rather than what the player may be able to do.

Objective messages: players need mission-relevant goal messages to understand what tasks they should complete, such as "defeat the boss monster in this dungeon."

Tutorial messages: to learn about how to play the game (to encode control messages and decode feedback messages), players will be sent tutorial messages.

System messages: system messages are necessary to remind players that they are playing a game on a computer that performs a variety of tasks, such as when the game is writing to disk, or that a log of activities has been updated in a menu. These are external to the game world.

This classification of interactions allows us to view video game interactions on a higher level for application to teleoperation. For example, if a teleoperation designer is designing an interface to communicate navigation aids to guide a teleoperator, they can explore the encodings for navigation messages; to communicate a continuous robot state, such as operating in an energy-conserved state, a designer could look to certain property submessage encodings. Thus, by thinking of the core idea of what they are trying to communicate, a teleoperation interaction designer can look at video game encodings for messages similar to that idea.

Below I break down each of these message types in more detail, gives examples of the messages and example encodings for the message, and describes how the messages are related to robots.

Property messages describe aspects of things in the game world, such as the player avatar, items in the world, or other characters. Example properties include health, energy, item quantity, etc. For example, a health message may be sent to describe player character health, the health of an enemy character, or the health of an item that can be broken.

There are two types of property messages: resources and states. Resources are aspects of a game entity that have a changing value, such as remaining ammunition [Halo: Combat Evolved, Bungie, 2001; Perfect Dark, Rare, 2000], health bars [Destiny 2, Bungie, 2017; Super Mario Galaxy Nintendo, 2007; Mass Effect 2, BioWare, 2010], or quantity of an item left [The Legend of Zelda: Ocarina of Time, Nintendo, 1998; Kingdom Hearts Series, Square Enix, 2002-2019; Overwatch, Blizzard, 2016] which can change if lowered or replenished. These messages are important for players so they can plan to accomplish their goal within the limits imposed by such dynamic resource levels.

State messages describe an ongoing state that is applied to a game object or not, such as an ongoing affliction [e.g., poisoned, see the Pokemon (Series), Game Freak, 1996-2019, Figure 6.8], or a beneficial state [e.g., able to travel 20% faster, such as running in Mass Effect 2, BioWare, 2010]. More complicated states that affect many controls and feedback interfaces at are also prominently communicated to the player. These *mode* states can include which actions a character can perform in that state [Metroid Prime, Retro Studios, 2002; Soul Calibur 2, Namco Bandai, 2002], or how a character will react to movement commands [Goldeneye 007, Rare, 1997; Uncharted 2, Naughty Dog, 2009; The Legend of Zelda: Breath of the Wild, Nintendo, 2017]. State messages are conveyed to the player as they describe a constant change to gameplay – a poisoned character (Figure 6.8) will have trouble in battle, or a stealth mode may have changed a character's movement speed. Both types of status messages can use similar encodings, or be encoded in the same interface, as they both describe quantifiable and discrete qualities about a character or object.

Application to Teleoperation: Robots also have multiple properties that operators may need to

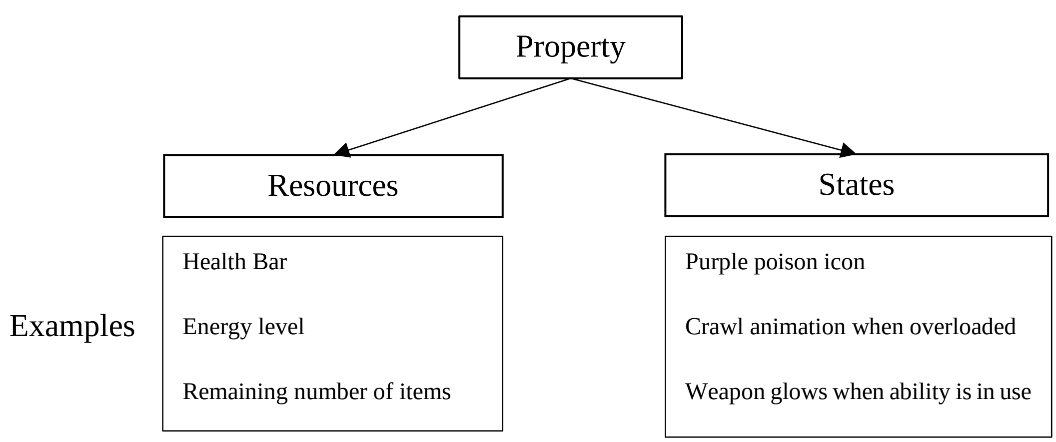

Figure 6.7. A breakdown of property messages. This is further broken down into consumable Resources, and continuous States. Bottom boxes are example encoding methods.

stay aware of: robots have resources, such as remaining battery level, or wireless sensors they can deploy; robots can also leverage state messages, such as being on slippery ground, having poor connection, or having broken end effectors. Similar to games, state messages can help explain non-standard behavior such as the robot not travelling as quickly as normal (the over encumbered state on slippery ground, or why a robotic arm cannot grasp an object like it may normally. Further, sensors can measure properties of other objects around the robot, such as distance, temperature, and more, and these properties of the environment may also be sent to the operator.

Interactable are messages that indicate if something can be interacted with, how something can be interacted with, and what the result of that interaction was. This general set of problems is actually highlighted in the situation awareness literature (Endsley 1988; Endsley 2016) and is important in games due the complex environments that are mostly non interactable (e.g., Figure) – this makes it difficult to understand what the player can do in an environment.

Interactable messages can shape how a player understands their abilities and explores the solution space for their tasks; as players are controlling an unfamiliar avatar with unfamiliar abilities, they may not think to use those abilities in new situations they encounter. For example, a treasure chest may be decorative only and not able to be interacted with at all [Dragon Age: Inquisition, BioWare, 2014], and this is common for objects in games in general (e.g., only some rocks are interactable in Figure 6.6, left, only some objects have interesting information to be "scanned" in Figure 6.10). If something can be interacted with, it should be conveyed in some way, such as highlighting the person or object with an outline (e.g., Figure

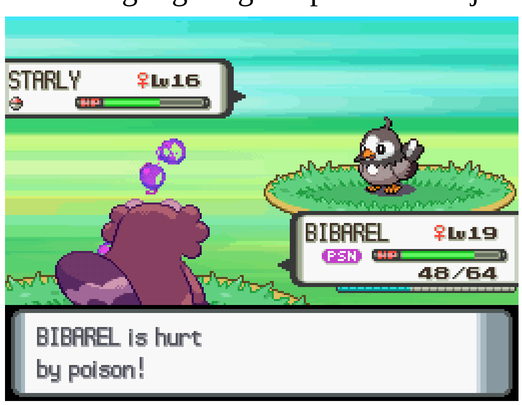

Figure 6.8. A player's character is poisoned. This message is transmitted with two property messages; one is the "PSN" (poison) purple icon near the player's health. The second is the animation that plays when the poison effect occurs and reduces health (a resource message – shown in the image as purple bubbles above the character). Pokemon Series, Game Freak, 1996-2019.

6.6, left). Knowing what cannot be interacted with also simplifies a player's mental processes: knowing what actions and items are at their disposal during a problem is important information to help a player formulate their next actions.

Once a player understands that they can interact with an object, a designer may choose to send a message about *how* a player can interact with that object. In the previous chest example, the chest may be openable by a player with a key, or (or in addition to that), be breakable to obtain the contents [Skyrim, Bethesda, 2011]. If the player tries to open the chest and finds it is locked, they may give up without thinking about breaking the chest. By communicating the breakability of an object, such as by displaying controls for attacking abilities when near a chest, the player may try different character abilities and actions to open the chest. This highlights the potential action space to a player when it may not be obvious.

A game interface can also help the player to predict the result of the interaction in some way, and this too may be conveyed in a message. By understanding the potential results of their actions, players can be better informed of their avatar's capabilities, and be aware of potential consequences for unfamiliar abilities. For example, a player may think of using a magical fire spell on the chest to break it, but this may result in the contents being damaged. This message may be conveyed when the player targets the chest with the spell. In other words, in addition to the range of possible actions, it can also be important to provide the user with an expected result if that actions is taken.

Finally, after the player decides their action, it is important to show the action, and convey the results of the interaction. This is important as the player can then evaluate whether their attempt was successful, and may contain useful information for how to proceed if it was not successful. Following the previous example, the player may have burned the chest, and the contents are theirs for the taking: the game should convey that the chest is now gone and the contents can be inspected and taken. Perhaps the chest is resistant to fire, and the attack does not break the chest: the game may convey this with a fizzling sound of the fire not working to convey that, not only was the attempt unsuccessful, but further attempts should likely not use fire as well.

Figure 6.10. In Metroid Prime (Retro Studios, 2002), multiple objects can be scanned for more information, but not all. Scannable (interactable) objects are given the orange icon to guide the player to understand what can and can't be interacted with.

Relation to robots: Robots typically have limited functionality (such as a small or weak gripper), and may not be able to interact with parts of the environment, such as opening a door handle. Communicating what could be interacted with, and what that result may be may help operators understand their capabilities: a system could communicate that an object is too wide to be picked up with a robot's effector, saving the operator time and frustration trying that solution, and they may instead try pushing it with the robot's body. Further, guidance on how to use the robot's tools, such as how an object may be successfully gripped, or where the operator should direct the robot to investigate may improve how well they can complete their goals.

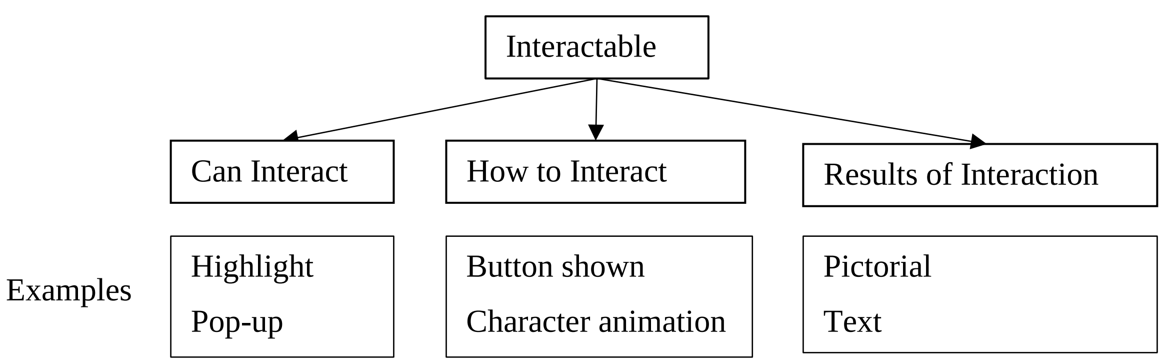

Figure 6.9. A breakdown of the Interactable archetypical message into sub-message types, and their common high-level encodings. These messages describe if and how a player can interact with a person or object in-game

Environment Awareness messages help the player orient themselves and events in the environment around them, understanding both where events are happening and when they will occur at that space in the world. Players are often moving in the virtual environment with other characters and important objects moving around them as well: a player may be fighting multiple enemies while escorting another character, who they must protect, to a location. In this example, the player needs to keep track of the enemies, which may be attacking, how they may be attacking, and if there is an opportunity to use any abilities against them. Thus, players may get disoriented, or forget about nearby characters when moving (Figure 6.14).

I emphasize that environment messages focus on space and time information for in-game characters, objects, and events. Spatial information may include location and orientation, of an enemy, or an attack [e.g., damage in Overwatch - Figure 6.12]. Events include time information, such as how close a bomb is to exploding [The Legend of Zelda: Wind Waker, Nintendo, 2002], when an attack will arrive in an area [Dark Souls Series, From Software, 2009-2016] or when an off-screen soccer ball will land in a location and can be hit [Rocket League, Psyonix, 2016, see Figure 6.13].

Relation to Robots: Teleoperators similarly have limited understanding of the remote environment, but can use sensors to overcome this. One could imagine detecting a voice or collision behind or to the side of the robot, the direction could be indicated such as in Figure

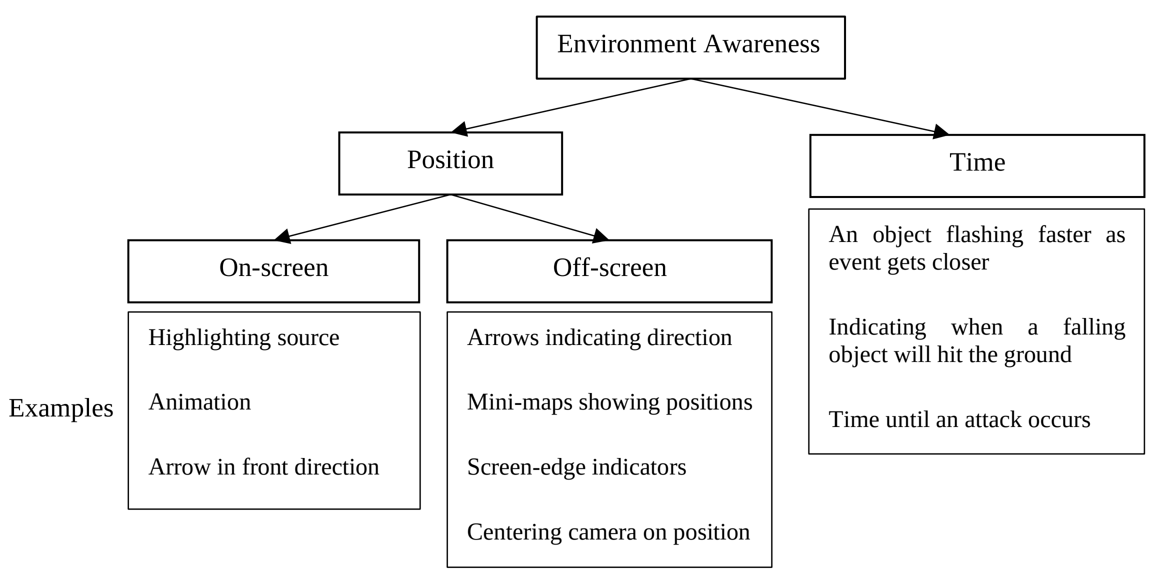

Figure 6.11. A breakdown of the environment messages and submessages: position, and time. Position is further broken down into on- and off-screen events. These messages describe where an object is or when an event will take place.

6.12. Other sensors may detect a person trying to pass from behind the robot, and indicate when they will catch up to the robot so the operator will not be surprised and can compensate for command lag over the network. A telepresence robot may be operated to follow a person, and if the robot turns away from this person an indicator could be displayed that communicates the person's position (e.g., to the robot's right) and distance from the robot. These messages

Figure 6.12. In Overwatch, damage taken is shown around the targeting reticule in the center of the screen. Note the small red wedges (highlighted in the white triangles). They are oriented exactly in the direction the damage came from. In the left image, the damage is coming from the player's rear, right hand side. In the right image, the damage is coming from the character in front of the player, slightly to the left. Overwatch, Blizzard, 2016.

Figure 6.13. In Rocket League, cars must hit a giant soccer ball that is often flying through the air. The game provides an environment awareness message that helps players understand where the ball is in the air around them. A virtual line is drawn perpendicular to the ground to indicate ground position. Two circles indicate how high the ball is – the closer the inner circle is to the outer circle, the closer the ball is to the ground. In the left image, this can be used to understand the ball is falling to the ground, helping the one car jump to catch the ball in the air (right). Rocket League, Psyonix, 2016.

can help operators form and maintain an accurate mental image of the remote environment, and react appropriately to events.

Navigation messages help players navigate around the environment or towards a destination. Rather than tell players where things are (environmental awareness messages), navigation messages tell players how to get to a position. Virtual worlds are often large and difficult to navigate without aids, much like the real world. Navigation messages can give general directions to a location [such as head east, or search in this general area, see Skyrim, Bethesda, 2011; Bioshock, 2K Games, 2007], the entrance way to the next area leading to the goal [Destiny 2, Bungie, 2017], even exact directions to take on streets [GTA IV, Rockstar Games, 2008, Figure 6.17] – see Figure 6.15. Further, navigation tools may simply present information to enable players to navigate on their own [e.g., mini maps in GTA IV, Rockstar, 2008; Metroid Prime, Nintendo, Retro Studios, 2002]. This map may further be decorated with environment awareness messages to better improve navigation (see above). I break down navigation message types in Figure 6.15.

Relation to Robots: Robots similarly need to navigate their environments. Operators may have specific destinations in areas that have existing maps that can be used to give navigation messages, such as a telepresence robot attending a meeting in a specific room. General

Figure 6.14. (left) the main character is trying to protect Queen Minnie from the enemies (left). The player must maintain environment awareness of the position of enemies and the queen, even while the camera rotates (right). If the player fails to understand this, enemies may get close to the Queen without the character noticing, endangering the mission. Kingdom Hearts Series, Square Enix, 2002-2019.

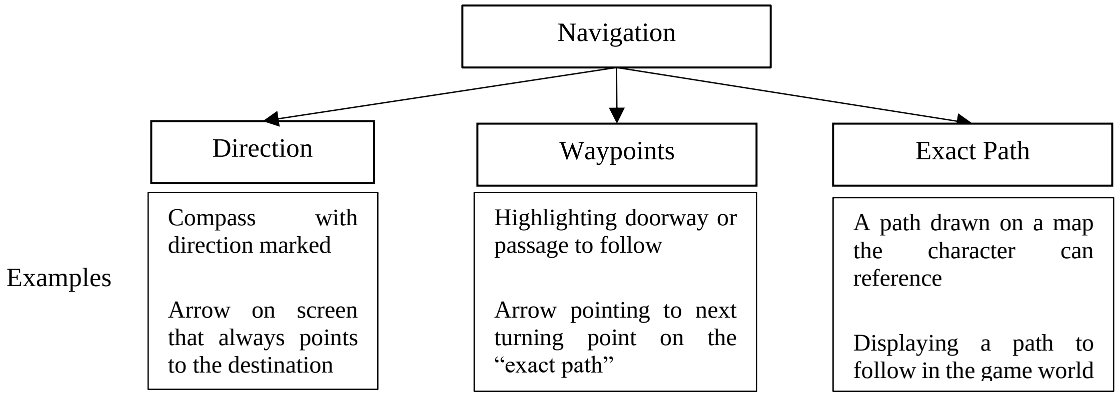

Figure 6.15. A breakdown of the Navigation messages that describe how players should travel to a destination.

direction messages can be used when the exact destination of an operator is not known, such as a search and rescue robot surveying an area. Further, games my limit navigation and map information explicitly to promote immersion or increase difficulty, such as by sending direction messages when the game could send an exact path. On the other hand, robots may not always have maps or full information about the remote area. Thus, robots may even leverage the limited-information interfaces (direction, waypoints) used by games when their own information is constrained.

Objective messages relate activity-level information to the player. Games generally have goals for players to complete. Games may even have multiple objectives, or objectives that change

Figure 6.16. Lines show the path respective teams (blue and red) to follow to reach the objective. This guides players in the right direction and fastest path, as well as helps them navigate to the other team's starting location.

as new information is discovered or an objective is completed. For example, a goal message may describe what is the next goal, how much progress has been made towards completing the goal, or any updates to those objectives (change in completion rate, new or changing objectives) – Figure 6.18. As each goal may have related information, important characters and locations, and even multiple steps to complete, games have devised encodings for helping players understand and keep track of these multifaceted goals.

Relation to Robots: Robot deployments are also often objective based. For example, a teleoperator may be using a telepresence robot to attend a meeting with a specific person at a specific place, a pipe inspection robot may have to inspect a given length of pipes for types of faults, or a search and rescue robot may have to search for survivors in a given area. These objectives may also change and update as new information is acquired. Conveying these messages to the operators can reduce the cognitive load of having to remember multiple objectives while operating. Thus, keeping track of all the objectives and what there is to do next can help teleoperators stay efficient.

System messages that detail things external to the game, but relative to the fact that the player is playing a game. While games often aim to be as immersive as they can, they sometimes need to communicate information that is linked to the system itself. This may indicate the game is

Figure 6.17. A zoomed in shot of the minimap in the bottom-left corner of GTA IV's interface. An exact path to the destination is highlighted on the map to aid the player's navigation in an unfamiliar world.

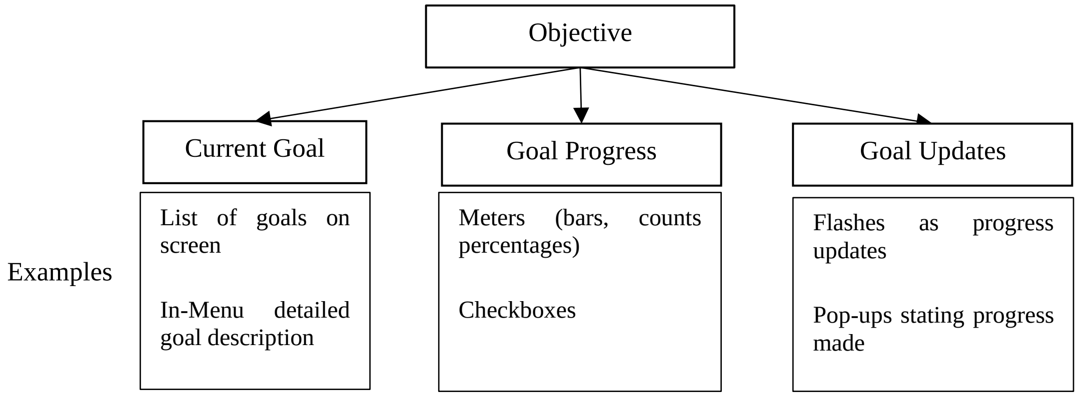

Figure 6.18. A breakdown of Objective messages that communicate what a player should be aiming to complete, and how close they are to completing that goal.

writing to disk and that the player should be careful of turning the game off. Alternatively, an online game may have a visualization of network latency to help players understand why they me see asynchronous behavior in their game. Communicating this to the player can explain erratic behavior, keep them aware of why certain performance lags may be happening, or inform them of other behavior of the system that may be useful to them.

Relation to Robots: Robots often suffer from problems due to hardware, such as a video connection that is unstable or encountering encoding or buffering issues, resulting in a non-smooth video connection that makes it hard to understand and react to the remote environment. Alternatively, the system may support logging, which the operator may activate or view reports on if the recording was logged succesffully. Communicating these states that are not specific to the robot or remote environment itself may help teleoperators understand and diagnose problems during teleoperation.

Tutorial messages explain how to perform actions and tools within the game, how to use existing actions in new ways, or help players when it seems they are having difficulty. Games employ these messages to teach players how to play the game without having to read instruction manuals. Essentially, they teach while playing. For example, the first time a player is climbing on a ledge, a tutorial message may convey which buttons to press to climb along the ledge [e.g., Batman: Arkham City, Rocksteady, 2011, see Figure 6.19; The Last of Us, Naughty Dog, 2013;The Legend of Zelda: Ocarina of Time, Nintendo, 1998; Metroid Prime, Retro, 2002]. This message may not appear a second time, or it may only appear in a reduced capacity, for example, just indicating the action is possible, as the ability to do the roll has already been taught. Well-designed tutorial messages can help a player grow to be more effective, give confidence, and help them utilize and understand all their character's abilities.

Relation to Robots: Tutorial messages may help remind operators of the robot's abilities (e.g., unused or useful sensors), help teach or refresh how to use a robot's ability (e.g., to use an end effector with a suction ability), or even suggest helpful actions in the current context (e.g., noticing if the robot is near obstacles and suggesting to activate or move cameras so the operator may better see the problem). Tutorial messages may be especially useful for beginner operators, and help those operators gain mastery of the system.

6.5.2 Archetypical Control Messages

Control messages are sent by the player to the system to affect the system state in some way – the player is controlling parts of the game. In spite of the range of actions I observed, my analysis revealed that controls typically fall into only two types of broad messages.

Figure 6.19. A video game has a pop-up text message detailing how to move in a new situation and what controls to use to complete goals in a tutorial message. Batman: Arkham City, Rocksteady, 2011.

Avatar control messages are the commands the player uses to make in-game characters and objects they control move or perform actions. While characters may be the player's avatar, or team members, avatar control messages may also change controlled objects such as vehicles, the cameras that allow players to see parts of the virtual world, cities the player manages, etc.

Configuration control messages interact with artificial configuration systems that affect how the game or avatar behaves. Configuration messages are an important part of game message sets as they can change how other encoded messages are decoded.
It should be noted that the design of control messages defines everything a player can do in a game. In general, the encodings are defined beforehand by the system designers: the move message may be encoded by pushing a joystick, and the player chooses when to invoke that message.

Avatar messages cover the majority of control design: how to make the avatar(s) in game perform actions that help the player complete their goals. My definition of avatar is broad: any abstract object in the game world that a player can control can be an avatar. For example, a player may send messages to make a character perform a jump action, or to make a vehicle accelerate forwards, or to rotate a third-person camera around the player. I describe two types of avatar control submessages: movement, actions, and novelty. *Movement* messages attempt to send one or more avatars to a new location. *Action* messages cover all things an avatar can normally do in the world. As these submessages are, themselves, broad, I treat each of them below as if they were their own archetypical message.

Movement messages from the player direct something in game to move. Movement messages may instruct an avatar to move from one position to another (almost every game I observed), reorient a vehicle [Mass Effect 2, BioWare, 2012; GTA IV, Rockstar, 2008], or perform a complex motion such as aiming a bow [Legend of Zelda: Breath of the Wild, Nintendo, 2017]. my analysis found that video games often use a small set of standard movement controls despite a player controlling a wide range of avatars, implying that those movement controls are useful and understandable for a variety of applications. My results further noted a common trend of more abstract, higher-level controls. This helps enable the aforementioned controls shared across a variety of controlled avatars, but also helps simply the player's workload. For

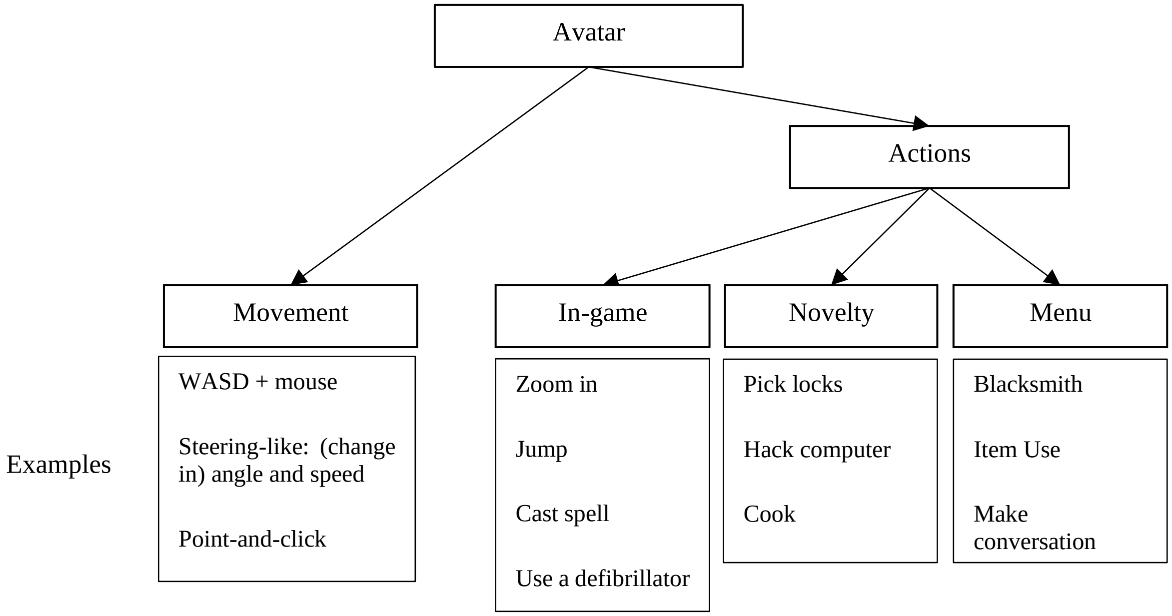

Figure 6.20. A breakdown of the avatar control messages and submessages.

example, while running forward, the player doesn't have to control each leg or worry about how to steer a particular vehicle, and games typically stop the player or disable damage from running into obstacles [e.g., Pokemon Series, Game Freak, 1996-2019; Final Fantasy Series, Square Enix, 1987 – 2019; Halo: Combat Evolved, Bungie, 2001]. In addition, this also makes skills gained in one game are transferable to a new game.

Relation to robots: Robots also have abstracted movement methods (such as the NaoQi API that can tell NAO robot to "move forward" without directing each leg joint. Further, robots can also have complex limbs and end effectors for difficult manipulation tasks. Controlling all of these in a fluid manner remains a problem in teleoperation. Video games' typical much higher level of abstraction could be beneficial here, focusing more on high level movement ideas that are shared across robots. Robot movements may even have a level of intelligence built in, such as automatically avoiding obstacles, in order to reduce the amount of tedious and detail-oriented work the operator must perform.

Actions are typically separate systems from movement, and are considered in the context of interacting with, or affecting the world. Picking which messages need to be sent to enable these actions is an important part of design. For example, a single button could have a character perform an entire fishing action: prepping equipment, casting the line, waiting, hooking the fish, and reeling it in the Pokemon Series [Game Freak, 1996-2019]. Other games take a slightly more manual approach, requiring players to send a begin fishing message, then send a reel-fish message when the player notices the fish-hooked feedback message [e.g., Final Fantasy XIV, Square-Enix, 2013].

Encoding actions is an extremely non-trivial task. Once a set of messages has been decided, they must be mapped to a physical control set. Games often enable players to send high-level actions by pressing just one or two keys, such as dexterously challenging tasks like reloading a gun [Half-life 2, Valve, 2004, Mass Effect 2, BioWare, 2010, etc], skill-intensive tasks such as blacksmithing [Skyrim, Bethesda, 2011], or long and complicated series of actions such as stealing a car [GTA IV, Rockstar, 2008]. In a role-playing game such as Skyrim [Bethesda, 2011] or Dragon Age: Inquisition [BioWare, 2014], there are hundreds of possible actions for a player to do.

I break actions down into three submessage types: in-game, novelty, and menu-based. In-game actions are those that happen in the default environment of the game, or the interface that is used while a character is moving through the virtual world. In other words, they are how actions are normally conceived. These include combat actions, basic environment interactions such as starting conversation and using items or tools in the environment, or activating other common actions like zooming a camera. Novelty and menu-based actions use different interfaces for feedback and control entirely, and are typically less frequently accessed modes in the main game.

Novelty actions are commonly available when video games enter a special mode for a specific task, with a custom set of control and feedback messages for that mode are sometimes referred to as minigames. At times, players may send messages that control small portions of the character's movement in highly specialized tasks, such as subtle hand movements for lock picking. In the lock picking example, a lock is shown and the normal orientation encodings (mouse or joystick) enable turning the lockpick, while movement encodings (WASD or second joystick) now attempt to turn the lock. These control encodings are only ever used in the lockpicking specific mode [Skyrim, Bethesda, 2011, Figure 6.21]. Other novelty actions include performing squats in an adventure game [Final Fantasy Series, Square Enix, 1987-

2019], cooking [Kingdom Hearts Series, Square Enix, 2002-2019], or playing instruments [The Lengend of Zelda Series, Nintendo, 1986-2018].

Menu-based actions are in-world actions that a character take that are completely mostly controlled through a more traditional menu. Skyrim again provides a good example with their blacksmithing system (Figure 6.22). A character does not have to perform all the individual actions of smithing, nor is there a single "smith" button. Instead, the character must navigate menus and select blacksmith recipes for creating a specific item, as well as the materials needed to use it. Then, the blacksmith process is performed automatically for the player. Similar types of menu-based action systems are used for crafting in other games [Dragon Age: Inquisition, BioWare, 2014], using healing items [Final Fantasy Series, Square Enix, 1987-2018], cooking [The Legend of Zelda: Breath of the Wild, Nintendo, 2017], or even all combat in turn-based game [Pokemon Series, Game Freak, 1996-2019]. These systems seem to be used most often when the player would otherwise have to have controls that could select from large numbers of options, which traditional menus are known to do well.

Thus, a player may be able to invoke many different kinds of control messages, and the input methods for encoding may be limited, such as a controller with 4-12 buttons. Even for a

Figure 6.21. A special mode with special movement control encodings for a fine-grained control task of lockpicking in Skyrim, (Bethesda, 2011).

keyboard that may be able to map all actions to a given key, this may not be ideal as the player must remember which key does which action, and be able to use the control method quickly, without thinking about the controls. Thus, video games tend to use controllers, or a limited set of inputs (only a few keys on a keyboard), and switch what those inputs do – they make the encodings modal, or sensitive to some context. This encoding mode may be dictated by a user-selected mode (see Configuration messages below), or in-game context (see interaction design strategies in Section 6.8). If the number of control options are too large, or the controls too specific for a task, menu-based or novelty control designs may be more appropriate.

Relation to robots: robots can have similarly complicated actions. Robots can activate multiple sensors, use complex arms or end effectors, move rubble in search and rescue, inspect and repair complex equipment, or even produce social signals like body language or facial expressions. Correctly designing these controls is a challenge (e.g., telling a robot to grasp a doorknob as compared to telling each arm joint to move in such a way that the operator can then grasp the doorknob), but from the examples in my data listed here, I can say there is a trend in video games towards providing a single message to invoke complex, high-level actions with a limited or specialized control interface.

Figure 6.22. A menu-based action in Skyrim [Bethesda, 2011]. A player chooses to perform blacksmithing actions. These are done by approaching an area with the correct tools, which opens a menu. In this menu, the player can select the item they wish to craft and what materials they will use to craft it. After selecting these in the menu, a smithing animation plays (i.e., the smithing is performed autonomously). This differs from typical in-game actions.

Menu based actions may also be appropriate on a robot with specialty tools, such as a robot performing science experiments. If the experiments have different variables or methods the robot may employ, these may be performed in a menu-based interface, with the robot performing the action autonomously. While this type of interface is not as immediately obvious for teleoperation application, as robots become more complex and can use multiple tools to perform a variety of tasks, such interfaces may become more useful.

Novelty interfaces are also relevant for robotics. Robots can be made to be specialists, with tools and abilities linked to these specialties. For example, a search-and-rescue robot may have health diagnostic equipment, a telepresence robot may be designed to convey social signals, or an inspection robot may be outfitted with marking or repair tools. In these cases, instead of developing a general control set for all robot actions, specific interfaces may be created *just to perform that one task*. While this may create more design work to create this set of interactions, this may be preferable for complex or delicate work where typical robot functions such as moving, displaying status messages about the robot or environment, etc., may not be relevant for the task at that moment.

Configuration control messages do not control a character, but access and change the properties of an avatar, the interface, or other virtual components of the game system. These include system-level *system* messages that change the state of the system itself or the game world, and avatar-level *mode* messages that change what messages a control encoding sends. These messages fundamentally alter the game as they change how encodings and decodings work. For example, in Figure 6.24, I can see that the encoding for a message "play ocarina" in the left picture is the same encoding as the "use a bomb" message in the right picture. In the same sense, the system receives the same encoding (button press), but decodes it to a different message in each case. The symbols in the interface themselves are showing property messages that describe the current mode (configuration) of each encoding.

System control messages shift fundamental properties of the system. However, like meta feedback messages, they are fundamentally about acknowledging the game is itself, a game, running on hardware and is made of software. Examples include adjusting the framerate or graphics settings in games, enabling or disabling non-game features such as subtitles, language

Figure 6.24. In The Legend of Zelda Series and other games, it is common to have limited access to all action messages at once. In the top-right of the interface, we can see 3 orange buttons that correspond to buttons on the controller. By pressing the button, the character will begin a complicated action, such as playing an ocarina, or throwing a boomerang. Note that the sets of actions between these two screenshots are different. These actions can be changed with Configuration messages.

Figure 6.23. An equipment menu encoding to change which control messages are encoded by specific button presses. This menu is how the encodings in Figure 6.24 change the messages they send.

settings, and more.

Games have controls not typical to normal productivity software, have limited input methods (discussed earlier), and need to be reconfigured in tense real-time situations. Thus, traditional menus and similar techniques may not be optimal. Further, game menus typically aim to not break immersion, and make great effort to maintain the pace, aesthetics, and feel of the game itself [e.g., Baldur's Gate II, BioWare, 2000 – see Figure 6.26]. More still, some games cannot be paused, and have developed systems for quickly navigating menus or other configuration interfaces during fast-paced, real-time activities. For example, many games that allow your character to have multiple pieces of equipment will let you switch equipment in a set order by

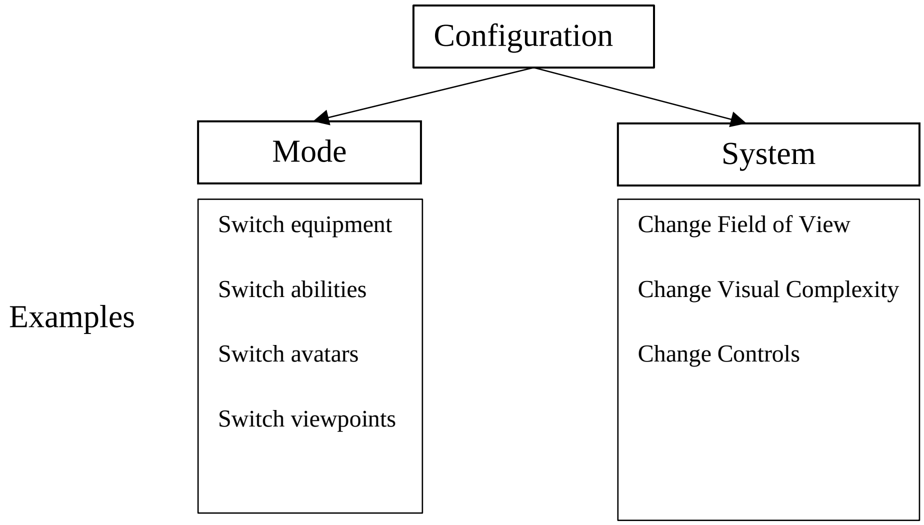

Figure 6.25. A breakdown of the Configuration messages. These enable players to change how they perform specific actions, which actions are available to them, or even aspects of the system the game is played on.

pressing a button or scrolling with the mouse wheel [e.g., Mass Effect 2, BioWare, 2010; Red Dead Redemption, Rockstar, 2010; Team Fortress 2, Valve, 2007].

Some configuration changes may not be messages explicitly sent by the player. For example, moving close to a friendly person may make the button used to swing a sword instead talk to that person [The Legend of Zelda Series, Nintendo, 1986-2018; Kingdom Hearts Series, Square-Enix, 2002-2019]. This is a contextual mode switch, and can be thought of as being implicit in the sense that the game is assuming the player does not want to attack.

Relation to robots: Robot interfaces commonly have complex menus for configuration, data storage, interface configuration, and more. Like some games, robots also do not have the luxury of pausing the world while menus are being navigated. Learning about how quick menus and in depth menus are used and structured in games could help make customizing a robot easier.

Software usability research has often cautioned about input modes as they can cause mode errors when the user incorrectly assumes which mode they are in (Norman 1983; Scarr et al. 2012). We, however, observed that many games have embraced modal input, potentially to keep physical controls simple. Complex controllers can make gameplay more difficult as a player has limited access to buttons with their fingers. If a game still requires dozens of actions and the controller only has eight to twelve buttons, modes can be used to give players access to all commands with their restricted controller. My data suggests that introducing modes while keeping controls simple may be a strategy that suits the real-time control problems of teleoperation.

Further, even game-centric system configuration ideas, such as changing how detailed graphics should be rendered could be used in teleoperation. In games, these adjust how much processing power is used by the game to add visual decoration, which increases processing power. In teleoperation, it may be useful to give operators similar control over what tasks the robot is spending more processing time on, increasing the resources available for other processes or saving battery life.

6.5.3 Summary of Messages

In these sections I provided an overview of the types of information games and players exchange with each other. This resulted in a surprisingly small set of high-level information types. I further gave suggestions to how each message type may be related to teleoperation. Interestingly, the messages I found were all compatible with teleoperation on some level. There may be elements of video game design that are less relevant, however, outside of interaction design, such as how the virtual worlds and levels are designed. Further, I selected my games to be genres that appear similar to teleoperation, which may explain why I did not find incompatibilities.

The applicability of core video game messages to teleoperation serves as evidence for my

Figure 6.26. Baldur's Gate 2 (left) and Destiny 2 (right) have both been praised for the design of their character configuration menus, both for functionality, and aesthetics (Candland 2016).

general approach of learning from video games to improve teleoperation. This was on an abstract level, however, and I did not yet discuss how those messages are implemented. The encoding of the messages are the actual user-facing designs and decide how users and the system interact mechanically. Thus, while considering messages helped me understand and categorize a variety of interactions, my analysis also covered the concrete techniques that realize them.

My message types are summarized in Figure 6.27.

6.6 Message Encodings

Throughout my analysis I saw numerous different message encodings. The design of a single encoding can be studied in great detail and have numerous effects on the user of a system (for example, Chapters 3-5). I would argue that such care is necessary to understand the unique challenges of each interaction technique: in my chapters I tested a number of different encodings for similar messages (e.g., make the operator focus on an area, or convey a sense of safety). Nearly all of my encodings were successful, and though they conveyed the same or similar messages, the differences between the encodings had sometimes subtle or surprising affects how they were interpreted by operators, and had different impacts on performance and

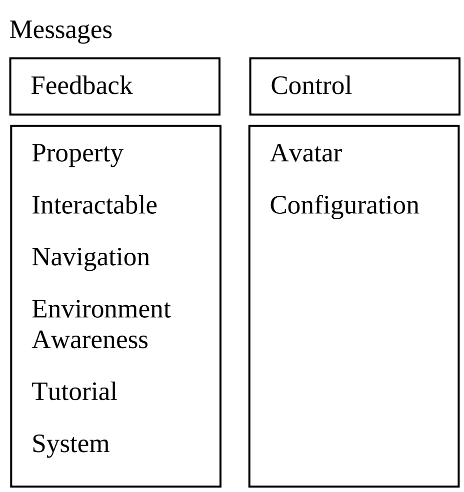

Figure 6.27. An overview of the message types discovered in our analysis.

experience. Encodings, then, should be investigated on a specific basis to see how it relates to teleoperation.

Understanding the impact of design decisions for encodings is further important as those decisions must consider how that encoded message be understood by the other party – decoded back into the original message. For example, to display player health in Figure 6.5, the number-only interface gives more precise detail, but may be decoded ambiguously by the player as it is unclear how the number relates to maximum health. In contrast, the bar interface gives less precise information, but may be understood (e.g., is the character almost full or almost no health) faster at a glance. In the case of control messages, a button press may be interpreted in different ways depending on system state; the system must send proper feedback messages to the user so the user stays aware of how their encodings of control messages will be interpreted (Figure 6.6). Understanding the relationship and impacts of each encoding and decoding in detail is complex, and may leverage aesthetics, design, psychology, social factors, the user's experiences, and perception (de Souza 1993).

One tool to understand these complexities emerged from my analysis of encodings: elements of encodings that were shared across all encodings I observed. They act as design components – decisions that must be made when making each individual implementation. I detail these *encoding parameters* below.

6.6.1 Encoding Parameters

I noticed common *encoding parameters*, or design decisions that appeared, in my sample, to be present in every interaction technique. While I argue that encodings should be considered holistically, understanding these variations that arise from changes in one or more parameters can better help me understand and explore the design space for a message encoding.

Every encoding had a design decision made for each encoding parameter. Parameters are independent from each other in that a value taken for one parameter does not change a value taken for another parameter. However, a choice of one parameter may restrict which values can be taken for other parameter.

Spatial: Encodings must have a position. A feedback encoding may be placed somewhere in on screen, embedded spatially into a 3D soundscape, or in the real world on a specific piece of hardware [like the use of a peripheral speaker in the controller in The Legend of Zelda: Twilight Princess, Nintendo, 2006]. For example, health bars can be placed anywhere in the overlay interface (Figure 6.5). Alternatively, the spatial component could link the physical and virtual worlds, such as with a force feedback controller for driving simulation games, with force feedback steering wheels simulating how a real car's steering wheel moves in different terrain. This hardware can encode haptic feedback spatially, such as rocks under the left side of your car.

Spatial components should be chosen to be relative to another object or absolute. For example, state property messages may be displayed relative to a character icon [Baldur's Gate 2, BioWare, 2000, Figure 6.28]. As a further example, a feedback message of a character's position may place the sound of the character's footsteps in the soundscape – a character behind the player character will be heard as if they are behind the player themselves [Overwarch, Blizzard, 2016].

Control spatial encodings are about where and in which space an action takes place. Similar to the feedback component, an action could take place relative to another object, such as a "swap to the next piece of equipment" configuration message– this depends on the current equipped item (e.g., rolling mouse wheel in Team Fortress 2 [Valve, 2007]). Alternatively, an action may take an absolute space, such as ordering a building to be built in a specific location in the world [Starcraft 2, Blizzard, 2010].

Relation to robots: Teleoperators should similarly think of the location of their encodings. For example, it has been shown that spatial relationships, such as how close things are placed

Figure 6.28. State property messages are shown relative to each character portrait, leveraging gestalt principles to associate each state icon with the character the state is applied to. Baldur's Gate 2, BioWare, 2000.

together on-screen, makes people perceive how related they are (Koffka 2013). Interfaces may further try to embed messages in camera feeds from the real world, or adjust or simulate a sound's position. Control encodings should make their spatiality clear; e.g., when a command is sent, it should be clear to the operator if that command is affecting a menu, the robot's movement, or its end effector.

Temporality: A message must be encoded at some time, lasting until another condition is reached, such as time passing or some game state being reached. For example, an environment awareness message may inform a player a defeated enemy is nearby (perhaps to be searched for a reward), but that defeated enemy may be hidden after some time – the encoded message ends (perhaps to reduce computational load on the system or visual cognitive load to the player, see [Resident Evil 4, Capcom, 2005]). An encoded character heath property message may only be shown when damage is taken and only disappear once full health is reached for some period [e.g., Super Mario Galaxy, Nintendo, 2007]. Alternatively, important messages may be encoded constantly, such as property messages that show which item is currently equipped (they are triggered on the configuration control message to equip the item, and do not stop until another item is equipped).

Control message temporality is about when the user's message is executed by the system. In other words, it is a question of command delay, or synchronicity. Some games allow command queueing, such as Starcraft 2 [Blizzard, 2010], where a string of waypoints can be entered for a controlled unit to follow once the last waypoint has been reached. Queues may have limits in length, where only so many commands can be entered before they are ignored [e.g. combat in Dark Souls Series, From Software, 2009-2016].This is contrasted with games that allow "button mashing" where multiples of a command can be entered and only when the current action is finished is a new command message allowed by the system [e.g., some attacks in Tekken 3, Namco, 1998; certain abilities in Overwatch, Blizzard, 2016]. Finally, sometimes commands can be interpreted immediately, allowing the system to be interrupted at any point, with the virtual character forced to cancel their current action and begin the new command. For example reloading a weapon in Destiny 2 [Bungie, 2017] can be interrupted by shooting the weapon, cancelling the reload, or giving a non-queued command in Starcraft 2 will overwrite and existing queue of commands. I note that this decision can be on a per-command basis. Fighting games [such as Tekken 3, Namco, 1998; Soul Calibur 2, Namco Bandai, 2002] commonly have make only certain moves "cancellable" while other actions must be seen through to the end, no matter the consequences. The key point of control temporality is it enables players to set actions that occur or continue into the future, as well as dictate when the commands end or can be ended.

Relation to robots: In teleoperation, feedback encodings could be context-sensitive, appearing when needed, and disappearing to decrease clutter in the interface, or reducing interface elements when the operator engages in a detailed manipulation task. Teleoperation designs should be careful of such triggers, as though it can reduce screen clutter it can also mean information an operator needs may not be available. While asynchronous controls have proven useful in teleoperation (Ashish Singh et al. 2013) control triggers could see increased use, especially with the time lag due to communication time and slow motors; borrowing timed feedback and asynchronous control methods from video games may make teleoperators better able to make use of their time.

Diageticness: A diagetic interface displays its information in a natural way in the environment. For example, instead of a typical on-screen ammunition counter for bullets remaining, a virtual in-game gun may have a readout on the in-game gun itself, listing the remaining bullets; this is a diagetic indicator as it exists logically in the game world. Or, energy levels may be shown by an actual battery meter of an in-game device [e.g., Overwatch, Blizzard, 2016; Destiny 2, Bungie, 2017]. I observed a continuum of diageticness: compare health meters in Figure 6.5 (page 199). The center image simply displays health as a number on a screen. A more diagetic interface is seen on the right, where the encoding is literally weaved into the virtual world's logic (also see on-character wounds in [Red Dead Redemption, Rockstar, 2010]; Figure 6.29). A semi-diagetic interface for health may be the health bars that are virtually overlaid on a person in the game world – this acknowledges the game world, while still being artificial from an in-game logic standpoint.

Control encodings have a similar continuum. A motion-sensitive controller may allow a player to swing a sword by performing a swinging motion in real life with motion controls [The Legend of Zelda: Twilight Princess, 2006]. Abstract representations of a motion may be used with a controller, such as moves low to the ground requiring the joystick also to be pressed down [e.g., Tekken 3, Namco, 1998]. Complex actions may also be activated with a single button press ["steal car" in GTA IV, Rockstar, 2008], which has virtually no diagetic aspect to it.

Relation to robots: Teleoperation has experimented with different levels of diageticness already. Robot paths can be drawn in the environment the robot will move in (D. Sakamoto et al. 2009), controls can be overlaid onto a video feed of the robot that can be manipulated with a touch-screen interface (Hashimoto et al. 2011), and non-diagetic interfaces have proven useful – tangible representations of the robot in the user's world (C. Guo, Young, and Sharlin 2009). Thus, considering what parts of an interaction should be diagetic is relevant to teleoperation, and video games provide new ways for robotics to leverage diagetics in design.

Granularity: While messages are abstract information, attention must be paid to the level of detail that the information is encoded. Navigation encodings are a common example of this – how to get to a location may be given by just a direction to head towards (very granular) [e.g., Skyrim, Bethesda, 2011], a series of waypoints (less granular) [Destiny2, Bungie, 2017], or a specific path (very granular) [GTA IV, Rockstar, 2008]. Resource encodings can also range in

Figure 6.29. Red Dead Redemption (Rockstar, 2010) does not use a traditional health meter, but shows increased amounts and darkening pools of blood as a diagetic health meter.

granularity, such as a specific numerical value that helps a player understand a resource vlue with precision, or a bar that enables quick estimation of that value, sacrificing precision for speed of intepretation (see health bars in Figure 6.5). How much detail is given can affect the mental effort required to process that information by the user.

For control messages, granularity describes how autonomous the action is. For example, aiming equipment in 3D space may be done automatically [Batman targeting the nearest grappling hook spot and attaching the hook perfectly with a single button press in Batman: Arkham City, Rocksteady, 2011]. Alternatively, aiming the item use may be completely up to the player [e.g. aiming a bow in The Legend of Zelda: Breath of the Wild, Nintendo, 2017]. In between these extremes is a range of "auto-aim" implementations, such as how a healing staff may work on a teammate as long as it is pointed roughly in the correct direction [Overwatch, Blizzard, 2016]. Control granularity can affect how challenging a task is, or how much agency a player may feel for that task. Too much granularity may alternatively make a repeated task seem arduous or tiring.

Relation to robots: Teleoperation designers may wish to consider performing less granular actions: instead of controlling a robot joint by joint, or moving an end-effector to a location, interfaces may instead encode higher-level actions such as "grab object" or "follow person." Feedback encoding design should consider *how much* information the teleoperator truly needs to complete their task (how many seconds left of battery use compared to an indicator that just displays 10% increments). Granularity is concerned with complexity, and reducing unnecessary complexity in either controls or feedback can reduce an operator's mental load, and perhaps help them better reason about how they can complete their tasks.

Medium: Messages must be passed through a physical medium. For example, a warning that health is low may be shown onscreen (many games), the game could create a warning sound to draw player's attention [e.g., Super Mario Galaxy, Nintendo, 2007], or the controller may be vibrated. Many control messages are created with haptics, such as buttons, mice, and joysticks. While not used in any games in my sampling list, camera-based (visual) input is famously enabled by the Microsoft Kinect peripheral. Sound has also been used as input, such as games using the microphone on the Nintendo DS handheld, or the microphone on the Kinect camera [e.g., Mass Effect 3, BioWare, 2012].

Relation to robots: Teleoperation research has largely focused on the visual medium, however my results suggest great potential lie in the aural and haptic mediums as well. For example, video games modify sound in the game so that more relevant sounds appear louder [teammates attacks and abilities are quieter than the enemies' in Overwatch, Bungie, 2016; Destiny 2, Bungie, 2017], or add in sound effects on state changes like someone noticing your movement [Metal Gear Solid 2: Sons of Liberty, Konami, 2001]. Sounds and haptics have been used to make events *feel* different, adding sound effects like mysterious sounds to make the player feel uneasy [Bioshock, 2K Games, 2007], or using haptics to make an explosion feel more powerful [The Legend of Zelda: Wind Waker, Nintendo, 2002]. These mediums should be considered more in teleoperation as they enable new ways to interact, and change the embodiment of how I experience the interaction (Klemmer, Hartmann, and Takayama 2006; Dourish 2001).

Aesthetic: this is a category for stylized variants of common interface techniques. For example, a detailed map may be stylized as a digital computer simulation [Batman: Arkham City, Rocksteady, 2011], or an accurate 3D recreation of the whole world [Skyrim, Bethesda, 2011]. Aesthetics is not simply about photorealism in video games, but about having a cohesive visual style (D. Sakamoto 2015, Figure 6.26 on page 222), often linked to the game itself, such as pirate, ocean, and sailing designs used throughout The Legend of Zelda: Wind Waker [Nintendo, 2002]. It is important to consider aesthetics as it can contribute to immersion or user experience of a game, and even make users more engaged and willing to use the software (D. Sakamoto 2015; Candland 2016).

Relation to robots: Teleoperation has largely focused on utilitarian interfaces, using standard look-and-feels for software design, or not considering aesthetics at all. Visually appealing interfaces and well designed and informative animations can help improve the user experience and make it more appealing to operate the robot. Looking to video games, better aesthetics could be used to create an urgent atmosphere or calming mood depending on the needs of the teleoperation system. While not directly related to performance, research has shown that such ideas are critical to success in interaction design (Carrol and Thoma 1988; Forlizzi and Battarbee 2004; Hochleitner et al. 2015).

6.6.2 Encoding Summary

My six encoding parameters teach me that concrete interface implementations have a number of shared dimensions that can be experimented with, discussed, and explored. As an example, Figure 6.5, page 199 demonstrates three different encodings for the same feedback message – player health. One is a bar that is filled proportionally to the player's health. The middle case is the absolute number of health a player has. The right side is another bar, but it is incorporated into the game world itself (the fluid along the spine of a special suit). Each conveys a variant of same fundamental message (player health) but displays it differently, based on where the encoding is on-screen (*spatial* parameter), how realistically they are embedded in the game world (*diagetic* parameter), how precise the information is (*granular* parameter), and their aesthetic (*aesthetics* parameter). All of these parameters can impact how a user understands the encoded message and the user experience the encoding creates.

I could apply encoding techniques common in games to support teleoperation, with encoding parameters helping to explore different design possibilities. Perhaps a teleoperation interface is visually crowded and it is hard to find space for a new encoded message in the virtual overlay – a designer may try to instead embed the interface in a more diagetic manner by placing the information over the relevant object in the robot's video feed, move to another medium like sound which will not interfere as much with the busy visual interface, or reduce the granularity of the information being presented to the user to allow for visually simpler encodings. Thus, encoding parameters are useful to analyze and discuss video game interaction techniques, but can also help discover new design ideas by exploring different values for each component.

I summarize the encoding components in Figure 6.30.

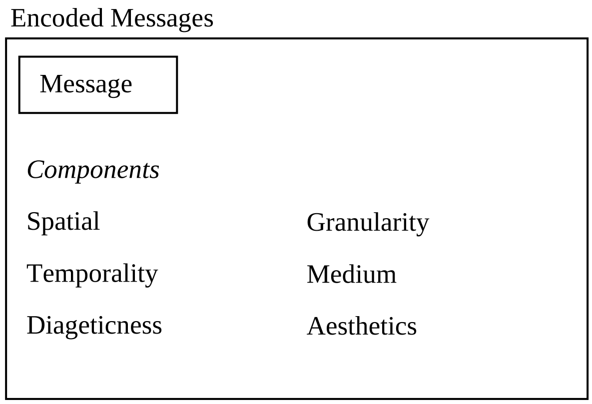

Figure 6.30. An overview of our encoding parameters

6.7 User Experience Goals

Messages are sent to convey information with encodings, which themselves have a complex design space. I developed these decisions – which messages to send, and how to encode them –alongside a set of codes about *purpose* in my analysis. In other words, I found common user experience goals that designs support across a wide range of messages, encodings, games, and genres. I noticed these themes were focusing on different aspects of the game experience, and they acted as guiding principles for how and why a message was sent or a certain design was chosen.

Arguably, most parts of a video game are design for entertainment in one form or another. Fun, however, is a complex and nebulous idea (Carrol and Thoma 1988) and thus in my analysis I strove to discover more descriptive types of entertainment in my data. This may be in the form of *satisfying* use of controls, designing *challenging* puzzles, or seeing an *immersive* world. From my observations, it seems that most aspects of interface design can impact the goal of fun.

But why should teleoperation designers care about fun? Prior work has argued that a good interface is not one that simply has no usability errors (Hassenzahl and Tractinsky 2006), and that non-functional experiences are critical in our understanding of interactions in general (Forlizzi and Battarbee 2004). Fun is one such part of the experience, and has been linked to increased motivation and engagement when completing tasks in software (Carrol and Thoma 1988; Deterding 2012; Li, Grossman, and Fitzmaurice 2012). Fun can create an atmosphere that changes how a person intrinsically engages with a situation, such as changing a feeling of uncertainty to curiosity, or motivating a person to learn something better (Malone 1982; Malone 1981). Thus, borrowing interaction designs from video games could alter a teleoperator's behaviour, help them improve at using a system, or just have a better experience and want to use the teleoperation system more.

I summarize my user experience goals in Figure 6.31.

Satisfaction: I noted design aspects that lead to a feeling of impact, *satisfaction* or *visceral* enjoyment. This theme was specifically related to actions, by either the system or player. For example, brief but noticeable pauses in animations as powerful attacks connect made actions feel more impactful [e.g., The Legend of Zelda: Breath of the Wild, Nintendo, 2017]; music and sound effects for narrative effect can make player actions feel heroic [e.g., Halo: Combat Evolved, Bungie, 2001]; or sounds and animation accompanying an in-game reward for finding important items, completing objectives, or winning a fight can feel even more rewarding (e.g., Super Mario Galaxy, Nintendo, 2007; Skyrim, Bethesda, 2011; Batman: Arkham City; Rocksteady, 2011; Final Fantasy Series, Square Enix, 1987-2018). The key idea is to make each action feel like it has impact in the world or on the outcome of a task, or to make those outcomes feel more significant.

Challenge: Challenge has a long history of being studied, and challenge's relationship with engagement is formalized in detail in Flow Theory (Csikszentmihalyi 1990). The goal is to provide a task with a difficulty that is not boring to players, while remaining possible to be overcome. As players improve, challenge can be increased with player ability, often seen in games with complexity layered on as the player progresses such as gaining new abilities in Metroid Prime [Retro, 2002]; The Legend of Zelda Series [Nintendo, 1986-2018]; Final Fantasy XIV [Square Enix, 2013], and more.

There is a genre of seemingly difficult games that defies this trend. The "Soulsborne" series [Dark Souls Series, From Software, 2009-2016] has recently become a genre in and of itself, defined by difficult combat, strong punishment for repeated mistakes, an immersive world, and little in-game help, aid, or tutorials. However, this genre actually follows the trend to flow or fun, as enemies in this game are designed with a limited number of movement and attack patterns that are *telegraphed* (giving cues and warnings for attacks and enemy movements to make them predictable and preventable with quick reactions). I see this fitting into flow theory in the sense that the game, at all times, provides opportunities to *improve*, if not necessarily overcome the challenge. Overcoming difficulty, or in this case, seemingly overwhelming difficulty can be enjoyable if designed correctly.

Immersion: Immersion is an oft-defined topic (Qin, Patrick Rau, and Salvendy 2009; Jennett et al. 2008; Mekler et al. 2014), but I use the term meaning feeling present in the game world. Games can be visually abstract, but I found general attention to detail, and the in-world consequences to player actions were related to immersion. For example, while riding a horse in Red Dead Redemption [Rockstar, 2010], controlling the horse's movement is reflected by showing your character's hands moving the reins correctly, rain drops bead and run down your character's helmet in Metroid Prime [Retro Studios, 2002], and characters react to the player characters actions [e.g., the crowd reacting in NFL2K, Sega, 1999]. In other words, the world *looks* real to the player.

This concept extends to include a sense of atmosphere, where video games try to evoke a feeling due to in-world events. This may be ambient sound effects in a horror game [Resident Evil 4, Capcom, 2005; Bioshock, 2K Games, 2007]. Computer character yelling threatening things at a player character to instill a sense of danger in Batman: Arkham City [Rocksteady Studios, 2011]. Games can create a sense of urgency in the player by adding swelling music and a countdown [Overwatch, Blizzard, 2016], or have computer characters die to add a sense of overcoming great difficulty, urgency, and consequence [Mass Effect 2, BioWare, 2010]. These atmospheric interactions create a *feel* of what it would be like to actually the exist in the game world.

Atmosphere is often used in the sense of our aesthetics encoding parameter. As a design goal, an atmostphere is a concept or high-level description of an experience. For example, a game like The Legend of Zelda: Wind Waker [Nintendo 2002] can be said to have a seafaring atmosphere, and this informs many parts of the designs: the music and sounds (e.g., waves, ocean breezes), the visuals (e.g., seagulls), and even the interactions. Instead of a zoom mode, the character uses an extending telescope; sailing becomes the main travel mechanic; maps gain an old paper aesthetic and navigation messages use compass designs, and an environment message of the wind direction gets sent so the player can adjust their ship. An atmosphere is the driving concept behind the game and its player experience and thus even interaction design decisions should try to comply with the atmosphere as best they can.

Understanding is a goal about helping players quickly understand the game world around their characters, and how they may interact with it. This combines a number of message types, such as property, navigation, environment awareness, interactable, and system messages, discussed in detail in the above sections. Note that this goal may be inverted by limited player understand to increase difficulty (see above).

Mastery: Another theme in my data was how messages that focused on instructing players (tutorial messages) or helping players master different control systems and explain the logic that makes the game function. While the presence of help menus is common in desktop software, video games often have tutorial modes for practice and teaching in situations similar to the real game. Some even have entire tutorial levels at the beginning of the game that slowly introduce basic functionality [e.g., Batman: Arkham City, Rocksteady, 2011; LittleBigPlanet, Sony, 2008]. Many video games will allow tutorials to be revisited, or even be customized to enable practice of difficult actions [e.g., Tekken 3, Namco, 1998; Soul Calibur 2, Namco Bandai, 2002]. To help players improve throughout the game, some elements are introduced to challenge players to play in certain ways or use certain abilities in new situations [DmC: Devil May Cry, Ninja Theory, 2013; Tony Hawk's Pro Skater 2, Neversoft, 2000]. I noticed that, like the challenge goal, tutorial tools should scale with the player to allow useful practice for different levels of mastery.

Player Agency and Automation: Related to designing difficulty, I noticed that control automation varied heavily between games for a given task. Automation can reduce difficulty, and many tasks in video games are performed even without player input. What I noticed was that games tended to make controls more manual the more those controls related to the core gameplay. For example, in shooting games, aiming is extremely manual, but other tasks such as picking up items or choosing what to say to other characters is either fully or mostly automatic. In contrast, some story-based games [e.g., Dragon Age: Inquisition, BioWare, 2014], a player may be given multiple choices in dialogues, with consequences and relationships to balance in those responses. Combat in this game has aiming all performed automatically, and only high-level commands are given to characters. Another example is having to compensate for wind direction or player skill when kicking in NFL2K [Sega, 1999] – it could be automated, but as it is part of the main game experience, designers make it more manual. Thus, even normally automated and menial tasks may be designed to take more thought or time from a player.

6.8 Design Strategies for Combining Encoded Messages

My analysis found strategies common across many games that help direct encoded message design to work together in an interface to support a design goal. For example, a character may have many possible commands, but the designer wants to have simplified controls for fluid interaction. They may then create groups of commands that are only usable in the appropriate situation, such as only enabling attack commands when enemies are around. I noted in my data that some combinations of encodings are archetypical across multiple video games. I describe the strategies I observed for encoding groups of techniques and provide exemplars of each strategy in detail, outlining their messages and encoding choices, as well as how they may be applied to teleoperation.

Managing Saliency: saliency is a term taken from visual information processing that describes how much something stands out visually to the human perception system. I borrow the core concept here to describe an encoding's ability to direct a player's attention generally, even if the medium is not visual.

If all feedback messages are sent at once, it is possible to overwhelm the operator; thus, games

User Experience Goals

Satisfaction	Understanding
Challenge	Mastery
Immersion	Player Agency and Automation
Aesthetics	

Figure 6.31. An overview of our framework of video game interaction with the addition of design goals.

have developed numerous methods to make high priority messages more salient than lower priority ones. One common technique to do this is to change the spatial positioning of a component to display important information where players are likely to be looking, such as damage markers in Overwatch [Blizzard, 2016], Figure 6.12 on page 209; or describing results of player actions near the objects or characters they affect (also where players are likely looking – Figure 6.32). This technique can only be used a limited number of times at once, due to the limited amount of space in the focal area but can be to help players better understand information critical to their immediate task.

The visual "camera feed" into the virtual environment is generally the highest priority in video games, and thus other player or environment properties are placed off to the side, making them less salient. These properties, such as health, may still be very important, however. Thus, updates to a health meter or ability can be animated (Figure 6.33), changing their visual design to draw attention. This use of animation to increase saliency and draw attention is common in video games, and includes targeting indicators (Chapter Three), damage indicators (Figure 6.12 on page 209), or to convey changes to player character properties like health (Figure 6.33), or actions (Figure 6.34).

Other modalities can make actions noticeable: games use loud sound effects to draw attention to an action result or game entity property. For example, if a designer wants to emphasize if an

Figure 6.32. When a spell is cast, the state properties are displayed around the enemy affected by that state (here, "Chilled" and "Frozen").

Figure 6.33. A healthbar has a static placement on screen. To draw attention to important changes, an animation is played with the characters face, as well as the health bar turning red for the portion of damage. This helps draw the user's attention to this interface. (Kingdom Hearts Series, Square Enix, 2002-2019).

attack hit or was blocked, they may use a loud or distinct sounds to better communicate this fact than visual animations alone [e.g., Soul Calibur 2, Namco Bandai, 2002], or to draw attention to a property message, such as emphasizing the character is standing in a dangerous area that will kill them [e.g., Metroid Prime, Retro Studios, 2002), or that the player has been seen by an enemy [Metal Gear Solid 2: Sons of Liberty, Konami, 2001]. Sound may further be altered to emphasize important things (nearby enemies, teammates calling for help), and reducing sounds from less important things (teammates that are not damaged, ambient noises) [Team Fortress 2, Valve, 2007; Destiny 2, Bungie, 2017; Overwatch, Blizzard, 2016]. Thus, video games commonly use sound to increase the saliency of a message's encoding.

Goals and teleoperation: salient interfaces can support many goals, including understanding (helping the player focus on the most relevant messages), satisfaction (emphasizing recent successful actions), or immersion (drawing attention to detail in the world). These techniques could help teleoperation in a similar number of broad ways. For example, a robot with multiple cameras or sensors could benefit from making only sensor data that has changed recently more salient. Immersive strategies could make the operator feel more embedded in the remote area, such as by providing direction indicators when an operator bumps into an object to better understand. Increasing the saliency of a successful action message may also improve the satisfaction of an operator, increasing their confidence in their actions.

While saliency and attention is known to be important to teleoperation (Chen, Barnes, and Harper-Sciarini 2011), I have found little work focusing on it. Some exceptions include my Chapter Three (Daniel J Rea, Seo, et al. 2017), understanding where operators are looking (Teng, Kuo, and Tara 2013), and using haptics to bring things to an operators attention (J. J.

Young, Tan, and Gray 2003). Video game inspired use of saliency in interface design may further help reduce cognitive load and improve operator performance.

Simplification is the strategy of compressing a large number of possible controls or interfaces to a simple set of encodings. This may result in very high-level and vague commands such as "Press Space to save Jack Ryder" to save a man from villains with a single button press [Batman: Arkham City, Rocksteady, 2011], or a player's ability to jump, run, and perform other feats of mobility may be represented by a single displayed statistic [Destiny 2, Bungie, 2017]. The granularity of the interface, as well as how they may be accessed (such as with different modalities, context-sensitive controls, or sequential and combination inputs, see Tekken 3 [Namco, 1998, Soul Calibur, Sega, 2002]) needs to be considered to maintain game flow and a feeling of control for the player.

Goals and teleoperation: simplification generally supports the understanding and mastery goals, as a simpler interface or control scheme is easier to learn and understand on the fly than a more complicated one. However, a simple design can also be an atmosphere or immersion goal as having simpler and fewer interfaces will allow a user to spend more time concentrating on the in-game world.

Teleoperation traditionally focuses on expert user interfaces that requires training and intense concentration. In recent years, a trend towards simpler interfaces has emerged (e.g., Double Roboics) for the same reasons and goals my analysis found games implement simplification. As video games are also often very complex and have found great success with the simplification strategy, teleoperation will likely benefit from encodings that support this strategy.

Context-sensitive interfaces are one common technique to reduce the need for complicated controls and interfaces in games. Video games often take into account the current game-state, or context, to change the interactions possible. For example, games such as Red Dead Redemption [Rockstar, 2010], or The Legend of Zelda Series [Nintendo, 1986-2018] have controls primarily based on exploration or combat, but when close to a computer-controlled character, the encodings used to attack or perform other aggressive actions can be used instead to talk or trade with the characters. This is a very common strategy in games, including controls normally used for making a character run automatically changing to cautious movement when moving along narrow cliffs [The Legend of Zelda: Wind Waker, Nintendo, 2002], a general "interact" button that can open doors, press buttons, pick up items, and more [e.g., Mass Effect 2, BioWare, 2010; Skyrim, Bethesda, 2011; Goldeneye 007, Rare, 1997; Uncharted 2, Naughty Dog, 2009; The Last of Us, Naughty Dog, 2013; see Figure 6.35]. These "modes" are induced by the environment, and cannot be changed except by changing that context in the game, such as moving the player character away from the area that is activating the context-sensitive mode.

Feedback messages may also be context sensitive. For example, environment awareness messages may change the camera angle automatically such that a player is always trying to move "up" on a top-down view of a sports field [NFL2K, Sega, 1999, Figure 6.36], or to maintain a good view of the situation [The Legend of Zelda; Ocarina of Time, Nintendo, 1998], or so that not everything can be seen in a stealth game [Metal Gear Solid 2: Sons of Liberty, Konami, 2001]. If actions are context-sensitive, such as in the above paragraph's examples, the context is similar to a configuration message mode switch (above), and this context or mode switch is useful to be shown to the player. For example, the sword icon shown near the button that swings a sword will turn into "talk" in The Legend of Zelda, Ocarina of Time [Nintendo, 1998], or how screen-overlay interface elements will become more transparent if they get in the way of in-world events [GTA IV, Rockstar, 2008].

Figure 6.34. A player's abilities are shown in the top-left corner. The right image shows a close-up. Some abilities take time to be able to be used again (a "cooldown"). To show time until the cooldown ends, the icon representing the ability slowly fills up, increasing interface saliency and helping stay aware of the incoming event of the ability becoming ready again. Mass Effect 3. BioWare. 2012.

Goals and teleoperation: Context-sensitive controls help support the control simplification goal, hiding actions that are only relevant in specific situations until those situations arise. Such modal interfaces are known to cause errors in software in general (Norman 1983), but can improve performance as users become more experienced, especially if proper feedback is given to the user (Scarr et al. 2011; Scarr et al. 2012). Teleoperation has mostly focused on lower-level control, enabling operators to solve problems manually at the cost of time. If combined with feedback messages to communicate what context the player character is in, teleoperation could benefit from this strategy as it can simplify the control space and present the most common actions for a given situation to the operator.

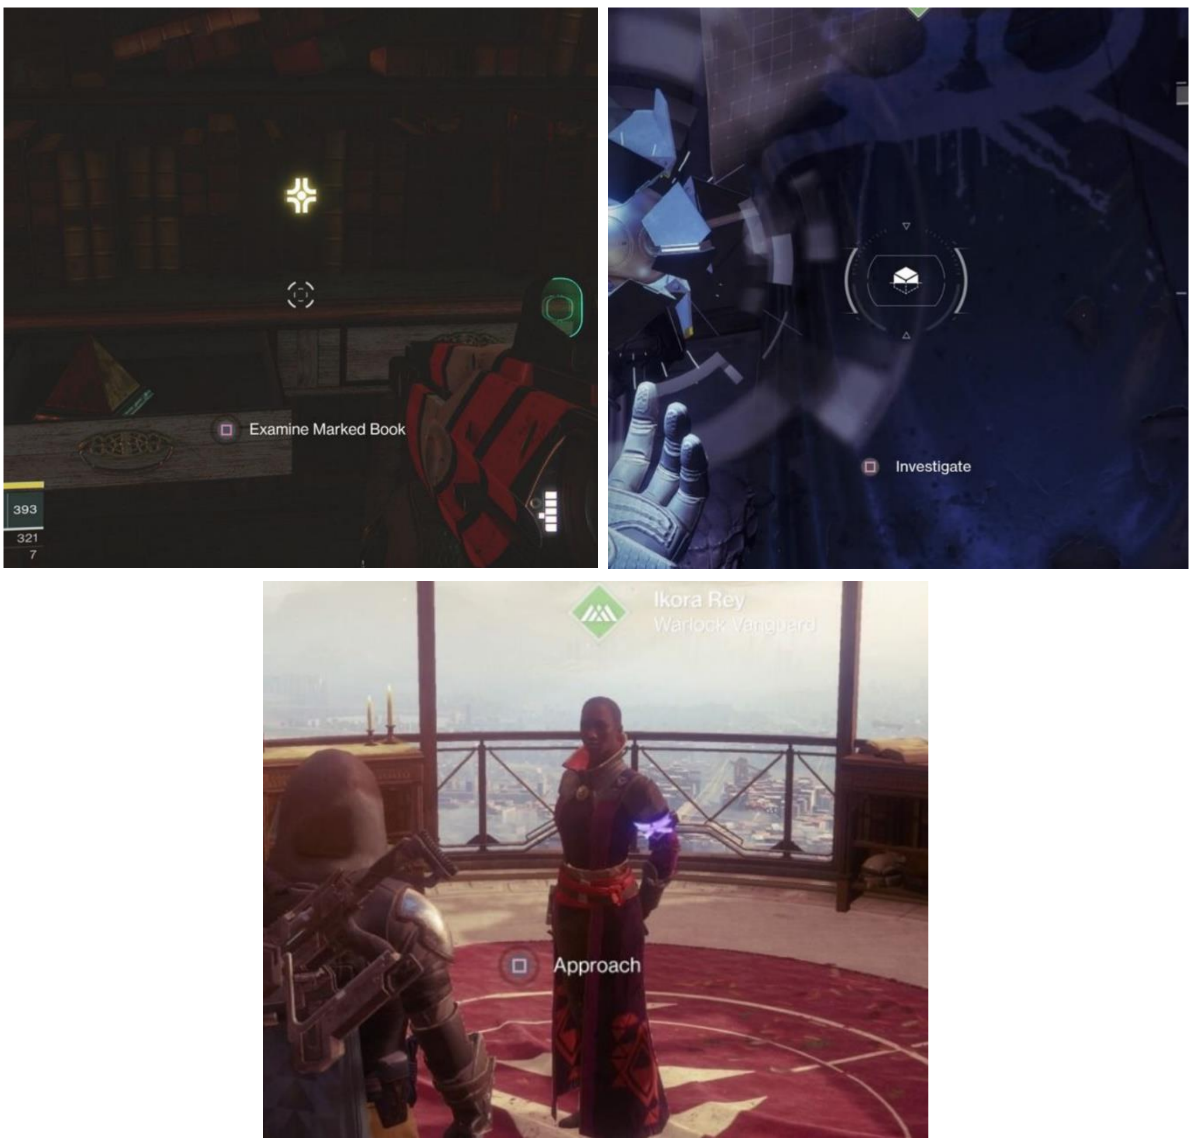

Figure 6.35. One common interaction technique is context-sensitive controls, where an avatar's environment is taken into effect to determine what action an encoding is used for. In these examples, one button (the square button) can encode for reading a book, investigating a symbol on the wall, or engaging a character in conversation. Destiny 2, Bungie, 2017.

Modal Interfaces are related to the idea of configuration messages and context-sensitive controls. Sometimes, performing one action may put a controlled character in a state that enables them to perform a new action or set of actions. In other words, a control message switched the control mode of the character. The key difference with context-sensitive controls is that the user actively chooses to enter a mode, where context-sensitive controls are activated by the system automatically through changes in the remote environment.

Fighting games do this frequently, such as how certain sequences of attacks must be entered first before finishing with a powerful attack in Soul Calibur 2 [Sega, 2002] or Tekken 3 [Namco, 1998]. Action games may also do this while fighting, with players selecting fighting styles on the fly, changing what their attack actions do [e.g., DmC: Devil May Cry, Ninja Theory, 2013; Kingdom Hearts Series, Square Enix, 2002-2019]. In these games, combination inputs may also occur, such as by combing action buttons with a direction to perform an action in that direction. In general, similar commands or sequential commands may be chained with similar buttons or button combinations.

Modes may similarly *replace* commands with others. For example, a game may let a player mode change from third person to first person to better see the character in their environment, or search and inspect respectively [e.g., The Legend of Zelda: Ocarina of Time, Nintendo, 1998, Figure 6.23, page 220; Super Mario Galaxy, Nintendo, 2007; Destiny 2, Bungie, 2017]. Other control encoding changes include targeting, where a user selects an object and all actions become relative to that target. For example, the generally global controls of forward, back, left,

Figure 6.36. In NFL2K [Sega, 1999], whichever team is controlled by the player, regardless of if they are on attack or defense, they are consistently moving "up" field. This is an example of context-sensitive feedback and gives a greater sense of consistency to players.

and right become “towards target, back away from target, strafe left around target, strafe right around target” [The Legend of Zelda series, Nintendo, 1998-2018; Red Dead Redemption, Rockstar, 2010; Kingdom Hearts series, Square Enix, 2002-2019]. Entire sets of commands or interfaces may be swapped with a button press, especially in games with large sets of controls and possible feedback like World of Warcraft [Blizzard, 2004], or Final Fantasy XIV [Square Enix, 2013], Figure 6.37. Like context-sensitive controls, it is important for the interface to reflect these changes to the user, resulting in modal feedback as well. The purpose of this technique is to enable the user to choose which tools they need to accomplish their goal in a quick and efficient manner.

Goals and Teleoperation: Modal inputs help the control simplification goal, and sometimes the mastery goal. If a person would like to have quick access to all buttons on a controller at once, it is easier to have fewer buttons, but fewer buttons restricts the number of actions that controller can perform at a given time. By using some buttons for mode switching, a user can

Figure 6.37. Numerous skills can be accessed by hitting the up-left-down-right or triangle-circle-x-square buttons on a Playstation 4 controller. If the player hits the ‘R1’ button, they can move to another set of commands (bottom image) that gives access to 16 new commands. This allows users to have sets of commands for specific tasks (battle, healing, etc). Final Fantasy XIV, Square Enix, 2013.

keep their hands near all controls at all times. For example, a robot may be able to be driven with a joystick, but also have a manipulator arm that can also be positioned to interact with the remote environment. Modal controls can allow a single joystick to be used for both movement of the robot and for moving the manipulator by switching into a manipulator mode, without having to move the operator's hands to another joystick. This is especially true if manipulation and movement are not expected to occur simultaneously. A search-and-rescue robot may have a set of tools to move rubble, and a set of tools to perform medical inspection of injured people they find; by creating a mode for each and allowing a user to swap between modes, a design can simplify the complexity of the robot control hardware by not needing dedicated controls for all actions in each mode.

Virtual Augmented Reality is a design technique used often in games, and is a medium-diagetic form of interface that artificially changes or augments information into the game world. For example, the quickest path to an objective can be seen in Overwatch [Blizzard, 2016] to understand where opponents may be emerging from, or how a kick direction and strength in a sports game may be shown by an arrow on the field [NFL2K, Sega, 1999]. This can be also used to cover technical limitations, such as areas outside of a renderable distance (based on computation limits) that are shrouded in fog [Half-life 2, Valve, 2004], or a game could show the path and landing area of a thrown object [The Last of Us, Naughty Dog, 2013, Figure 6.38].

In addition to overlaying information virtually, a system can also provide feedback by distorting the game world itself. For example, games sometimes add in screen shake to emphasize an important event [Baldur's Gate II, BioWare, 2000; Destiny 2, Bungie, 2017]. Virtual weather effects like fog can be added to indicate unknown areas or outdated information [Starcraft 2, Blizzard, 2010]. Music in a 3D soundscape, emanating from an area the player should head to could act as a navigation message based on the audio modality [Bioshock, 2K Games, 2007]. Sound could be virtually edited in real time to emphasize sounds originating from dangerous areas [Overwatch, Blizzard, 2016; Team Fortress 2, Valve, 2007]. Or, physics could be virtually broken, like projecting a shadow under a character jumping, to help understand where they will land, regardless of the environment lighting [Super Mario Galaxy, Nintendo, 2007, Figure 6.39].

Goals and teleoperation: this strategy can support the understanding goal. Teleoperation can benefit from many of these techniques directly. For example, if a building map is known, the path to a meeting room can be virtually displayed on the floor in reality from the view of a telepresence robot operator. Environmental awareness of a teleoperator can be improved by highlighting objects and increasing saliency for an item that may be of interest (a potential person that needs help or an object the robot knows how to interact with). An interface could enhance events, like making an object that is in the robot's path glow to warn the operator, or overlay the name and other information of a person above their head during conversation.

Teleoperation could use the concept of virtual augmented reality to embed more information in a natural-like way, perhaps increasing immersion and enabling people to utilize their own intuition of physical reasoning to understand the remote world. For example, a fog at a distance may be used to indicate the limits of sensors; shadows could be virtually cast on the ground directly under objects to help operators better understand the spatial relations in the environment; adding in screen shake on a small collision to make the collision seem more dangerous and encourage the operator to drive more carefully. Virtually-embedded controls (showing controls on a third-person view of the robot) have already seen to be successful and help reduce mistakes and improve performance (e.g., Hashimoto et al. 2011). Virtual augmented reality can be as a compromise between the needs for a flexible interface that also

Figure 6.38. The system simulates the character's throw to show the path and landing point using the augmented virtual reality strategy. [The Last of Us, Naughty Dog, 2013]

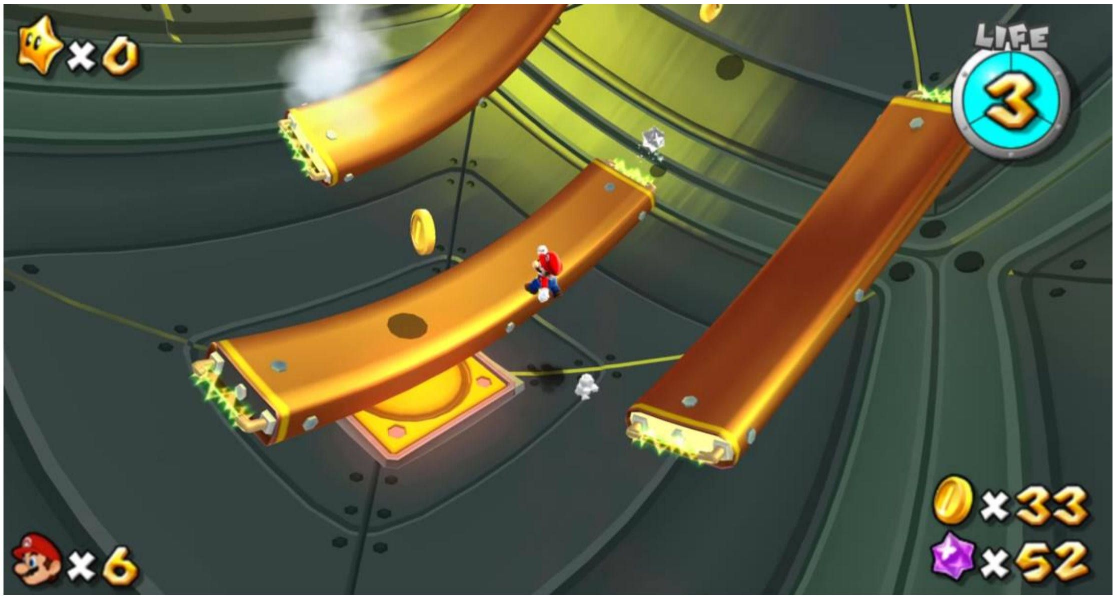

Figure 6.39. The shadows are manipulated to always appear right under the objects (see the coins, Mario). This enables better understanding of verticality and 3D position in a game that requires precise jumping. [Super Mario Galaxy, Nintendo, 2007].

moves toward immersion by leveraging the operator's knowledge of the remote world.

Standard Control Encodings: while a tenant of this thesis has been that video games may be a source of innovation, video games themselves leverage common and expected control and feedback formats from other games. For example, many games have adopted common movement controls. For controllers, this consists of one analog stick being used for translational movement, and one for rotational movement. Computer versions of these games will have translational movement performed by the "WASD" method, with keys being used to move forward (W), back (A), strafe left (A) and right (D), while orienting the player character with the mouse. These movement controls were almost universal in the games I observed that had first or third person camera views [Mass Effect 2, BioWare, 2010; GTA IV, Rockstar Games, 2008; Bioshock, 2K Games, 2007; Half-life 2, Valve, 2004; Spec-Ops: The Line, Yager Development, 2012; DmC: Devil May Cry, Ninja Theory, 2013, Dark Souls Series, From Software, 2009-2016, and more]. Alternatively, top-down camera views all relied on asynchronous controls, with mouse clicks setting the positions of a move action. These top-down games further shared concepts like scaling the granularity of commands – grouping multiple avatars together and giving them an identical command [Baldur's Gate II, BioWare, 2000; Starcraft2, Blizzard, 2010]. Interestingly, one game allowed the player to mode switch between a top-down and third-person camera; each camera view also switched to the related control scheme mentioned above [Dragon Age: Inquisition, BioWare, 2014]. This provides further evidence for how certain presentation styles and genres use similar basic movement controls.

Standard controls are often shared within genres, with the above examples roughly belonging to the “First-Person Shooter” and “Action Adventure” genres. Looking to other genres, racing games had their own standard controls, which used a forward-facing camera with similar controls: one button for acceleration, one for braking, one for checking rear-view mirrors, etc. [Mario Kart 7, Nintendo, 2011; Gran Turismo, Sony, 1997]. Interestingly, within each genre the games appear quite different: driving comic go-karts compared to driving sports cars. Or in the examples in the above paragraph, players are space soldiers battling aliens, assassins in ancient history that ride horses and sneak around cities, or scientists driving vehicles and climbing through caves trying to uncover mysteries. Yet, control messages are shared between them, and thus they are able to use similar encodings.

Goals and teleoperation: Leveraging standard interfaces can help the goal of Mastery, as players can take skills from one game and apply them to others. I can imagine this could be leveraged in teleoperation by developing standard sets of controls based on task or robot type. For example, telepresence robots are used to explore a space and help the operator talk to people and are often a simple tablet on a movement platform. If these platforms used similar control schemes, an operator could log in to a new robot, and still be able to perform their tasks well immediately. Similarly, robots with arms or other manipulators could share common interactions for common manipulator tasks, such as directing a robot to pick up an object.

Combining like messages and encodings is one common strategy that ends up creating interfaces with many encodings compounded into one element. A common feedback related example of status bars. I observed a common encoding of “danger:” a threshold of a remaining resource for a player character that triggers other messages to bring the player’s attention to the situation. When a resource, such as health, is low, a warning sound can be played while this state continues [e.g., Super Mario Galaxy, Nintendo, 2007; Overwatch, Blizzard, 2016; The Legend of Zelda Series, Nintendo, 1985-2018], and the bar could denote the remaining health in some way, such as by delimiting it [e.g., Destiny 2, Bungie, 2017] or changing the colour of the resource indicator [e.g., Pokemon Series, Game Freak, 1996-2019]. These health bars often have other health-related state information, such as poison (Figure 6.8 on page 204), extra assets like shields and armor [Dragon Age: Inquisition, BioWare, 2014], or other commonly used resources by combining different resource bars together [Final Fantasy XIV, Square Enix, 2013, Figure 6.37; Kingdom Hearts Series, Square Enix, 2002-2019, Figure 6.33].

Goals and teleoperation: combining like messages is a strategy that can improve understanding, and grant an opportunity to consider aesthetics as a means to combine the encodings in a coherent manner. By keeping similar encodings together, it will be easier for a user to process all the information at once, as they will be relevant in the same context, or use similar decoding processes on those encodings, reducing their cognitive load. Teleoperation can leverage these ideas as well, better organizing sensor information and robot state in the interface. Further, this pattern suggests teleoperation should consider how to change related information to similar encodings, perhaps by experimenting with encoding components until an understandable and aesthetic design has been achieved.

Temporal Interfaces: to further limit the amount of information shown at one time, video games leverage encodings with different triggers to create interface designs that only exist for a given period of time, and then change or disappear. For example, health bars in video games are often not shown when the character is at full health [Super Mario Galaxy, Nintendo, 2007; Mass Effect 2, BioWare, 2010; Uncharted 2, Naughty Dog, 2009], or only when health changes [Goldeneye 007, Rare, 1997]. Objective messages may only be encoded when an objective changes or is completed [Gran Turismo, Sony, 1997; The Legend of Zelda: Breath of the Wild, Nintendo, 2017]. Status information may be shown for a few seconds after a controllable character is selected [NFL2K, Sega, 1999]. The key idea is to give players the information they need only as they need it and remove it when it becomes less necessary.

Controls may also take time to encode. One common idea is that a button must be pressed for a given period of time to, for example, drink a healing item [The Last of Us, Naughty Dog, 2013]. Alternatively, modes may only be entered after an action is performed for a period of time: in The Legend of Zelda Series [Nintendo, 1986-2018] and the Kingdom Hearts Series [Square Enix, 2002-2019], making a character run into a large and heavy object for a short time will make them enter a "pushing mode" in order to move the object. These timing delays help imply that there will be consequences to cancelling an action, (e.g., dropping the healing item), or make the controls simulate the effort of taking a long drink or pushing an object.

Goals and teleoperation: time sensitive interfaces can help increase understanding by providing timely information, improve aesthetics by reducing screen clutter, make actions more satisfying with interface elements or controls that make an action seem more impactful. They can increase immersion by making actions that take time in the game take time for the player to enter, and change a task's difficulty by requiring use of a control message to be held for a time. This is a broad area for application, and teleoperation could similarly benefit from considering *when* an encoded message should be triggered, and what should trigger that message to disappear. The goal of control simplification could also be furthered in teleoperation by actions that require longer button presses to begin or finish. For both feedback and control, temporal interfaces can help simplify the interface.

I add design strategies into my final framework in Figure 6.40.

6.9 A Framework of Video Game Interaction Design

My analysis identified a framework that can explain all the techniques I observed in terms of four key components. My analysis method resulted in a set of **archetypical messages** sent between the system and user to support interaction. These messages are encoded into concrete implementations whose designs vary broadly. Nevertheless, I observed that certain **encoding components**, design decisions that must be made when encoding any message in my results.

When choosing the messages to support the interaction or designing an encoding, **user experience goals** can be used to anchor decisions made. These goals can be supported by **interaction strategies**, which are general approaches to designing and combining concrete encodings. I note that the concrete techniques I presented are a small set of examples I observed within my dataset. The breadth in video game interface design extends even further.

I combine all the components of my framework and their relationships to one another in Figure 6.41.

6.10 Should We Design Teleoperation Interaction like a Video Game?

This chapter leveraged a message-passing conceptualization of interaction to dive into 853 interaction observations of 30 games. What I found was a surprisingly small set of categories in each framework level that explained most of my observations. Further, my analysis found emergent higher-level strategies and design goals developers may use to drive their decision making. This provided a layered view of video game interaction design, allowing me to compare each dimension of my framework to teleoperation in each section.

Overall, I found that many of my observed messages, encodings, goals, and strategies were relevant to teleoperation, and provide an avenue for learning from video games to improve teleoperation. In contrast, I know of few specific techniques in teleoperation leveraging game design; thus, while the general categories of interaction appear applicable to teleoperation, more research is needed to investigate the effects of specific interaction designs.

Our work and framework suggests that video game interaction design is broadly applicable to teleoperation. Even with my sample of video games, I generated numerous interaction techniques that may be adapted to teleoperation, and there are still many more games and game genres to investigate. Indeed, I purposefully limited my sample to teleoperation-like video games, but games that do not share obvious similarities may hold techniques that can inspire

Design Strategies
Interface Saliency
Context-sensitive interfaces
Modal Interfaces
Virtual Augmented Reality
Standard Control Encodings
Combining like Encodings
Temporal Interfaces
Simplification

Figure 6.40. A summary of our design strategies.

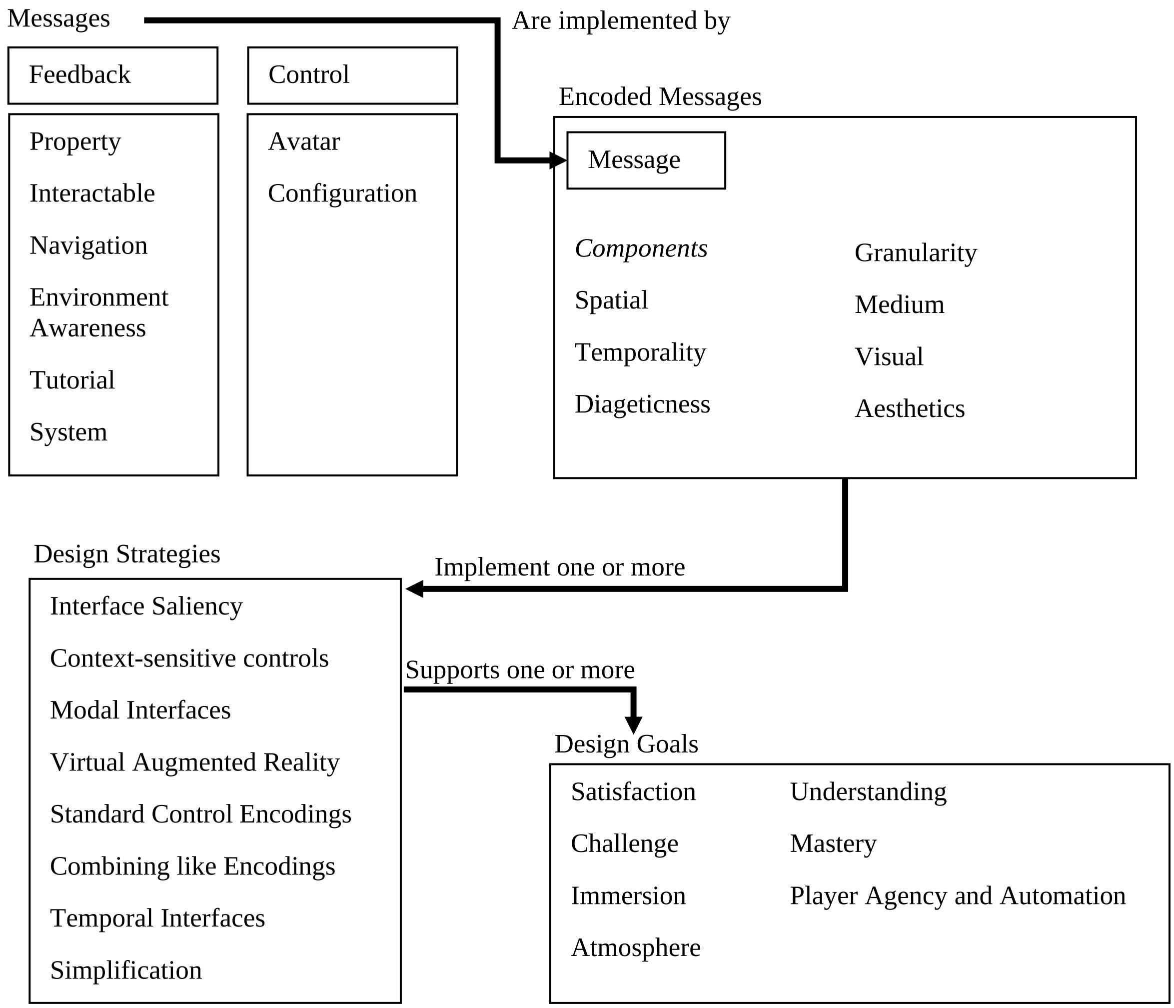

Figure 6.41. An overview of our framework of video game interaction and the relationships between each of its components.

radically different or innovative teleoperation designs – even if those designs will take more thought and research to adapt to teleoperation.

It is interesting to think that I did not find large categories of video game techniques that are not completely useless to teleoperation design. While some ideas, such as state property messages, interface saliency, or certain control schemes are clearly applicable to teleoperation, others were not. For example, social interactions with computer-controlled characters does not seem necessarily useful in teleoperation, but taking the *idea* of leveraging our social natures to impact operator behaviour proved to be useful (Chapter Five). Others may seem physically impossible, such as abilities that allow player characters to see through walls. However, even this idea, on a higher level, is about fusing other sensory information (say, x-ray vision or echolocation) into the main video stream, which is an idea actively under research in teleoperation. It therefore appears that some techniques need to be abstracted to better understand their purpose and their effects before being adapted to teleoperation. In other words, we must not just understand *what* video game interaction techniques do, but *why* they were designed in that way, and what their effects may be even on a more subtle and unconscious level. I thus recommend deep and explorative research when design and evaluating video game inspired interaction designs for teleoperation.

In particular, even in the video games I observed, it is difficult to make note of all design choices and interactions. Some interaction ideas, if done well, are difficult to notice, such as techniques with modified sound levels, or augmented reality that changes the position of shadows on-screen: our brain perceives and translates these into meaning without effort. Game developers use these techniques to, in some sense, lie and trick players into having certain experiences or understanding certain information without the player noticing. My Chapter Four work on priming in teleoperation touched on these types of interactions, which I only noticed to game developer's admissions to the deception (Scheurle 2017). These revelations suggest that directly working with or observing game designers and their rationales and playtests may reveal more of these designs made to go unnoticed.

My work presents a large space of future work to explore in teleoperation. With the evidence in this thesis and this chapter, I have demonstrated underlying theoretical similarities of the two fields, and how the overarching strategy of learning from video games can benefit teleoperation. I leave the broadest discussion to Chapter Seven, but believe my work on the framework is a seed that can be extended and broadened by interaction work in teleoperation in general.

New teleoperation designs can be compared to those found in my framework, and if games in it share similar designs, it can help predict the effects of those interfaces while further providing evidence for my approach. Exploratory research into video games by themselves can also expand my framework; more games from more genres can be studied, or those chosen with different sampling techniques could discover new messages types, encodings and components, goals, or design strategies. Such data could add further depth to the framework, for example by discovering new sub-messages. I especially see importance in studying the effects of the different dimensions in my framework, which can help better researchers better understand the relationship between user experience, interaction design, fun, and performance.

6.10.1 Reflection on Message Passing for Analyzing Video Game Interactions

Throughout my observations, I noticed that during interaction, the many parallel messages, encodings, and decodings quickly become complex (Figure): both the user and the system may be sending multiple messages simultaneously (visually, aurally, and through haptics), some messages and responses may be asynchronous while others may be synchronous, and some interfaces may dynamically compound encodings for multiple messages. However, even in such cases, the message passing paradigm provided me with tools: I was able to break interface elements down into their individual messages, analyze how that interface element encoded a message, and reason about how it is interpreted, or decoded. Thus, message passing has a scalability by describing complex interfaces via the layering of multiple messages, while understanding encodings and decodings can help design how to communicate the messages effectively between the user and system.

6.11 Conclusion

The work in this chapter provides the community with a broader look at the interaction designs used in video games; I abstracted the interaction problems being solved in game interfaces using a message-passing paradigm, and organized the interaction design goals and techniques of those interfaces into a framework. I used the framework to demonstrate how video game interaction designs solve similar high-level problems also faced in teleoperation, and my framework highlighted connections between video games and teleoperation that can be leveraged for novel teleoperation design. Modern human-robot interaction designs were shown to be similar to other designs in my framework, serving as further examples that my approach is generalizable. The classified video game interaction techniques from my survey serve as examples for future inspiration of teleoperation interaction designs. In general, my framework provides the vocabulary and structure to discuss video game interaction at a higher level for applicability to teleoperation, as well as provides evidence and theoretical grounding for the broader applicability of my approach.

CHAPTER SEVEN: DISCUSSION AND RECOMMENDATIONS

Across the projects in this thesis, I saw video game interaction techniques applied to different teleoperation problems. I have explored interface design that directs operator attention, how the presentation of a robot can alter task performance and perception of a robot's capabilities, and how social agents can be used to influence operator emotion. I further developed a framework that looked to what video game interfaces aim to do and how they do it on an abstract level, identifying common interaction designs and approaches across a broad range of game interaction techniques which I used to discuss the similarities between video games and teleoperation..

I briefly summarize the lessons learned in each chapter:

1) Chapter Three: games direct attention – inspired by video game damage and targeting interfaces, I designed, prototyped and evaluated several new interfaces grounded in perceptual psychology that help direct operator attention to areas on-screen. I further outlined design guidelines based on these game-inspired interfaces for future designers.
2) Chapter Four: games manipulate players' perceptions – one trend in video games is to aim to alter the perceptions of players to create specific experiences. I applied

this idea to teleoperation by priming operators about the physical capabilities of robots. I found that I could affect an operator's perception of speed and safety of a robot, and change how safely they drive, even without changing the robot itself or its performance.

3) Chapter Five: games use social designs to make players feel – games often create social interactions in a game that can create empathy or relationships with the player. I took this idea and created a virtual passenger that reacts to the operator's driving safety. My results showed that we can influence an operator's emotions based on how the agent reacted, demonstrating that social interaction designs could be used to improve teleoperation.
4) Chapter Six: what game interfaces do and how it relates to teleoperation – I explored the similarities between teleoperation and video game interaction design more broadly by building a high-level framework that describes and organizes video game interactions from multiple angles. Specifically, I discussed the core information that video games and players exchange, how they exchange it, video game design goals, general implementation strategies, and specific interface techniques. I examined how each of these angles relates to common teleoperation problems and solutions. Linking these components together, my framework provides one approach for understanding and discussing video game interaction design and their broader use in future teleoperation designs.

Reflecting on these four angles together, I can draw lessons spanning across the thesis. My goal was to explore how learning from video game interaction techniques could benefit teleoperation as a general approach; from the evidence in these projects, I learned that video game inspired interfaces can shape a teleoperator's experience to improve task safety and efficiency, as well as change the way the teleoperator views their own performance and the qualities and abilities of the robot. From these findings I make three broad recommendations:

1) Improve User Experience to Improve Teleoperation Performance – I found that my video game-inspired designs improved user experience. Further, I found that altering

user experience could also lead to improved teleoperation performance, though the two are not necessarily linked.

2) Shape User Experience by Designing for Human Psychology – I shaped operator experience by targeting different aspects of human psychology including visual perception, priming and expectations, and social psychology. Thus, we can improve teleoperation experience and performance by designing for how people react and think, in addition to the traditional approach of focusing on functionality of the interface or robot.

3) Good User Experience is Generally Beneficial in Teleoperation – I found numerous benefits in my results for targeting user experience in my designs outside of just performance increases. We could improve (reduce) self-reported operator workload, operator emotional state, as well as influence how they perceived their performance and the robot's capabilities.

In the remainder of this section I detail these recommendations and discuss my findings.

7.1 Improved User Experiences Improve Teleoperation Performance

Game designs take into account a variety of cognitive, social, and emotional psychological processes in their designs. Many interface designs such as status displays, navigation aids, and environmental awareness interfaces are designed to reduce the player's mental workload and help them solve in-game problems. In my survey, I frequently found that game interfaces aim beyond pure functional design; many interactions are designed to have an aesthetic, to make a person feel a certain way, or to produce an emotional response in the player. In other words, they aim to create a specific user experience, in addition to a goal of communicating information.

My results demonstrate that designing for user experience in teleoperation is a key approach for impacting operator driving performance. In particular, considering an operator's behaviour and thoughts about the interaction can improve and inform designs in ways that are not obvious from simple performance data. For example, in Chapter Three, for a visual search task I had similarly performing designs in terms of how many targets participants found, but some designs created a higher workload. This extra dimension relating to the operator's experience, not simply their task performance, helped me understand that some designs will make it more difficult for the operator to maintain their performance over the long term (see Chapter Two: Section 2.1.2 on cognitive load; Sheridan and Simpson 1979; Parasuraman, Sheridan, and Wickens 2008).

I additionally influenced people to perceive their performance differently: in my priming results, some of my robot presentations garnered strongly-worded feedback claiming the robot, in comparison to the other robots, was frustrating to control, was too fast to be safe, and more, despite there being no actual difference between the robots participants drove. Even when I did modify the robot response and driving performance, I found an impact on the operator's self-reported driving feel or experience, making them feel less rushed or better performing regardless of their actual task performance. In general, I found that design changes that affected the user experience contributed to or correlated with improving the operator's driving behaviour. In this thesis, I pioneered the approach of focusing on teleoperation user experience to improve performance.

Thus, I recommend that user experience should be considered a core and integral component of teleoperation system design. When the concept of user experience was entering human-computer interaction design, Hassenzahl and Tractinsky suggested the field should focus on creating "outstanding user experiences rather than merely preventing usability problems." My results agree with this conclusion (refer to Chapter Two, Section 2.1.6 on user experience), and go further: I found that, in teleoperation, user experience *itself* is a way to improve usability, performance, and an operator's perceptions of themselves and the system. From my observations, video game designers embrace this: games strive to be fun, engaging, immersive, and satisfying to use. Video game interaction design, therefore, can act as springboard to bring better user experience to teleoperation.

7.2 Shape User Experience by Designing for Human Psychology

Across all of my studies, a common theme was to be user-centered, and focus on different aspects of user psychology – cognitive processes, social psychology, etc. – to understand and influence how a user will feel and act. I have demonstrated the importance of moving beyond core usability and human factors in teleoperation design – ergonomics, interface layout, etc. – to include these components of how people think about, perceive, and react to different interaction designs.

I found video games use the approach of understanding different aspects of human psychology to better design their interactions. In my framework, I described how video games have design goals that may aim for certain psychological effects such as fun, understanding, and satisfaction. Other research has also noted that the leveraging of these types of mental processes is a common difference between video games and normal computer software (Langer, Hancock, and Scott 2015). These designs affect the player's user experience, which may impact user performance, perception, and more. Thus, video game interaction design is a source of new designs and inspiration for leveraging these effects in teleoperation.

7.3 User Experience is Intrinsically Valuable in Teleoperation

In the prior section I discussed how my results link improvements to usability and operator performance to improvements in user experience. Here, I argue that – irrespective of performance gains – user experience itself is valuable. Research already links good user experience to improved technology use, adoption, and use patterns (Bargh, Chen, and Burrows 1996; Lindgaard et al. 2006; Mitra and Golder 2006; J. E. Young et al. 2009; Hassenzahl and Tractinsky 2006), and my results extend that body of work by demonstrating that changes to user experience can affect a teleoperator's perception of their performance and the robot, as well as influence the operator's emotions.

While each study presented in this thesis had its own design goals and performance metrics, I found that I could purposefully design to shape teleoperation experience with different aspects of interaction design (perception, emotion, and even social qualities). In the remainder of this section, I outline effective strategies I found for designing for user experience with video game interaction designs throughout the thesis.

7.3.1 Changing Experience by Directing Operator Attention

To shape user experience, I can purposefully design interactions to influence where an operator looks, and what they focus on. In this thesis I learned of human physiology and how our visual and neural systems process information, and leveraged this to draw operator attention to areas to specific areas on-screen.

Designers can direct attention to improve overall operator experience by reducing cognitive load (Chapter Three) and by guiding attention to important interface elements (Chapter Five). In both cases, my goal was to limit the information operators process to the most relevant by increasing the visual saliency of specific areas. Conversely, if visual perception is not considered explicitly, a less important interface element's design may draw an operator's attention too much, becoming distracting, increasing cognitive load, and reducing long term performance. Thus, considering how a visual design (visual aesthetic, animation, etc.) will draw attention can help guide operators towards specific behaviours and experiences.

7.3.2 Influencing Operator Perception and Expectations of their Experience

My priming studies demonstrated the power of designing to affect an operator's conceptual perception of a robot (Chapter Four) – how various stimuli and experiences intertwine to influence how people think and behave. By priming operators about a robot's capabilities, I changed how they perceived their robot's weight, steering, safety, speed, and even sometimes made operators drive more safely. Thus, by shaping operator perceptions and expectations, designers can make people feel that their teleoperation experience is more dangerous, safer, more successful, etc., without adding to or modifying the functionality to the robot itself. My virtual passenger interface and driving profile priming similarly found a trend for affecting an operator's perception of their own performance. Instead of making the robot seem harder or easier to drive, we can also change the operator's *perception* of how they behaved, which is linked to operator performance (Hart and Staveland 1988; Steinfeld et al. 2006). Thus, we can design interactions to influence an operator's perception of an experience without modifying related parts of the robot or interface.

7.3.3 Influencing an Operator's Emotions During Teleoperation

I demonstrated how teleoperation interface design can shape operator emotions. In Chapter Five, I changed the positivity of an operator's emotion by the end of their task; in Chapter Three, my attention capturing designs could evoke feelings of urgency, calm, or stress, similar to how my driving profile priming made people more or less rushed. These changes in emotional state can influence how people behave when driving (Leshed et al. 2008; J.A. Groeger and Rothengatter 1998; John A. Groeger 2002), or when using software (C. Peter and Beale 2008; Langer, Hancock, and Scott 2015). If robot designers do not consider how their interaction design can affect the operator, they risk consumers forming incorrect and negative opinions as well as performing worse due to those feelings.

Taking into account an operator's emotion or feelings can provide further insight into design problems. For example, a potential customer may be nervous and scared of hitting people or obstacles if they are accidentally primed that a robot can drive very quickly, or an operator may feel disappointed and unconfident due to the way performance metrics are displayed. I found tools and methods from other fields could successfully measure changes in operator emotional state, such as by measuring emotions with standard questionnaires or simply analyzing operator's general qualitative responses about a teleoperation system for overall affective themes. Thus, new teleoperation work considering operator emotion can use these tools to gain insight into how an interaction affected an operator emotionally.

7.3.4 Influencing Operator Experience with Social Agents

I pioneered the idea of bringing social agents and the use of social techniques into the teleoperation interface by taking cues from computer-controlled characters and other social interfaces in video games. This evidence that techniques from social human-robot interaction can also be applied to teleoperation agents to affect the operator lays the groundwork to further leverage work in social human-robot interaction that can help with other teleoperation problems, such as communicating task or robot state (Admoni and Scassellati 2017; Breazeal et al. 2005; Sharma et al. 2013b), directing attention to parts of an interaction (Vázquez et al. 2014), and changing levels of trust in a robotic system (Banh et al. 2015; Hancock et al. 2011).

In addition to exploring additional techniques from social human-robot in a teleoperation context, further social techniques may be leveraged from video games. Social interaction with artificial agents appeared frequently in my survey, being used to draw attention to areas in-game, give recommendations of strategy to the player, or even make the player question their decisions from a moral perspective. Story-based games, such as those made by developer BioWare, extensively develop and employ social interaction game mechanics, and even specifically advertise how much you can interact with and even romance the in-game characters (Roberts and MacCallum-Stewart 2016). While this is an extreme example, design lessons and goals may be taken from works like this to increase social engagement and communication with social techniques in new teleoperation interaction designs.

7.4 Learn from Video Games to Improve Teleoperation

Together, my thesis lets me broadly recommend learning from video game interaction design to improve teleoperation. Video game interactions are designed to improve user experience by leveraging human psychology, and I applied this approach to teleoperation: designing, implementing, and evaluating a number of video game-inspired interactions. My results demonstrated that video game interaction design can improve teleoperation performance and experience.

In addition to my specific projects in my chapters, I used my framework to make broad comparisons between general classes and archetypes of problems and interactions in video game design to those in teleoperation. For example, I found that video games and teleoperation share many types of basic information to be displayed to a user (such as system state, position in environment, or how to navigate safely to a destination), or control needs. I further discussed the potential for video games design goals or implementations strategies to be useful in teleoperation, even though the goals of video games (to be entertained) may not always seem aligned with the task- and performance-oriented nature of teleoperation. Thus, teleoperation can learn not just from the lower-level information passing and implementation of video games, but also from video game goals and overarching design strategies.

The abstraction in my framework and my projects themselves also had the benefit of hinting at how some video game messages or interfaces could be unexpectedly leveraged in teleoperation. For example, it was not immediately obvious that social computer-controlled characters were useful in teleoperation. However, by understanding what those characters were conveying, such as tactical or moral advice, pointing out interesting areas in the environment, or calling out if they noticed traps, I realized that social agents were just another encoding method, conveying similar messages as more traditional interfaces. In this way, even video game interaction design can serve as a source of new and unexpected design approaches for improving teleoperation.

In these ways, my framework provides an abstract way to explore an interaction's goals, core message, or general design strategy in order to develop future teleoperation designs. By thinking of a problem in teleoperation, my framework can provide examples of how the interaction design would be approached in video games, serving as a starting point for a teleoperation design. This approach is validated by my project chapters, where I created interfaces that mitigate common teleoperation problems with methods used for similar effects in video games.

In short, I found video game interaction designs can improve teleoperator performance by improving user experience. I further found that improving user experience in teleoperation had many other benefits for the operator, such as reduced cognitive load, better emotional state, and improved perception of performance and robot ability. Video games often aim to produce these effects by understanding different parts of human psychology and designing an experience around them; these game design can also apply to teleoperation. Thus, user experience is an important component of teleoperation design, and leveraging video game interaction design as inspiration can serve as a springboard for new teleoperation designs that aim to improve the operator experience.

CHAPTER EIGHT: LIMITATIONS AND FUTURE WORK

In this chapter I outline directions and problems for future work. My discussion highlighted the importance of user experience and its effects on operator behaviour, however throughout this thesis I encountered limitations in quantifying both user experience and teleoperation behaviour. I discuss these methodological challenges as well as the limits of how I can generalize my results, both in terms of what this thesis has shown I can learn from video game interaction, and as well as how these lessons can apply to teleoperation at large.

In addition to improvements in methods, my work has implications beyond video game inspired interfaces in teleoperation. I detail two promising directions for teleoperation research outside of direct video game inspiration: the potential for user modelling in teleoperation, and the ethics of manipulating a user experience. These future work directions showcase how other areas of research interest can be generated from investigating video game interaction design.

8.1 Limitations

I encountered numerous difficulties throughout the chapters in this thesis. I classify these into two broad groups: the limits of the generalizability of my approach, and my experimental methods. While I explored video game-based teleoperation interaction designs both at a detailed level (Chapters Three-Five), and at a broader level (the framework in Chapter Six), I am limited in my ability to generalize to video games at large due to both my sample size, and sampling methods. My experimental methods were also limited: I encountered difficulties throughout this thesis to quantify "teleoperation behaviour" and "user experience." I reflect on the methods I did use, their implications for my results, and some suggestions for how to improve similar methods for future research.

8.1.1 Generalizing Video Game Interaction to Teleoperation

While I can confidently say that we can learn from video game interaction design to improve the user experience and performance in teleoperation, I must also state there are caveats to this conclusion. This generalization does not apply to all video game interfaces; in particular I limited my sample to only video games that already share characteristics with teleoperation. If I broaden my sample, such as to include rhythm, puzzle, or text-based games, I expect to encounter interfaces that may be radical for teleoperation and not obviously applicable, like specialized input hardware for music and rhythm games.

Additionally, my sample was limited in scope. Within my 30 games observed, I began to see considerable overlap in the interfaces and themes I recorded. While this implies that I was converging on common design patterns used across multiple games, I cannot say they generalize to games even within a genre. I was still encountering multiple new interaction techniques for each video game I observed; I would prefer my new observations to almost entirely fit into my existing framework or not generate new themes and codes. Due to limited resources I have yet to reach this point, and so the research is ongoing.

I chose critically acclaimed video games as a sampling method, I noticed similarities within the games in my sample. For example, many top-rated video games were made by the same studios (e.g., BioWare, Nintendo, etc.). As studios may have the same designers work on multiple projects, this will limit the diversity in my sample as the same people will be more likely to use previously successful designs, limiting the breadth of interactions I observe. Even within genres, I saw interaction conventions shared heavily across games. On one hand, this made it easier for me to identify prominent design techniques within those studios or genres, it once again limits me to commenting on looking to video games in general for future teleoperation interaction designs.

Thus, I can say that critically acclaimed video games within my definition of “a video game that is similar to teleoperation” are likely to be useful as inspiration for new and useful teleoperation designs, but I am limited in my ability to generalize beyond this without broadening the video game survey. My framework helped me gain an understanding of the variety of game techniques available and understand their general applicability to teleoperation. However, there are still many more games to investigate. Thus, one way to continue to expand on this work is to add to the framework by exploring the vast space of video games (Section 8.2.1).

8.1.2 Methods for Video-game Inspired Teleoperation Interfaces

In the previous chapter, I recommended attempting to design for and observe changes in the user experience as well as how the experience affects operator behaviour. One of the major limitations I had while attempting this myself was methods: I found myself limited in how to quantify changes in experience and measure teleoperation behaviour. I found limited resources in the teleoperation literature, and those I used and developed myself lacked power to provide descriptive results. Due to my strong prior recommendation of pursuing good user experience design in teleoperation, I believe this to be one of the most important problems to move this area of research forward. I discuss these struggles, and how I believe future work can improve upon us.

Measuring Experience and Behaviour

It is unclear what user experience factors are important for teleoperators. I outlined the background of user experience in human-computer interaction in my related work, however these more general metrics may not be sufficient to measure user experience in teleoperation as what is important about a user experience can be related to the application area. For example, user experience in video games can be evaluated by metrics that may not be applicable to other software due to experience being the main goal of the software as opposed to productivity (Hochleitner et al. 2015).

I found that general goals drawn from video games may provide improved teleoperation experience, and the applicability of these experiential measures need to be experimentally verified. Even further, good teleoperation user experience will likely have unique requirements unrelated to other software including video games. Thus, exploring and defining user experience guidelines and heuristics is an important step for improving teleoperation via user experience.

If I have targeted user experiences I want to create or improve, I will also need better tools and instruments to measure the outcomes of my designs. One of the key difficulties I encountered was choosing measurements to evaluate my designs. As user experience is broad, I chose measurements that similarly spanned a number of concepts that range from performance measures, measures of cognitive load, measures of emotional state, or changes in user perceptions about certain qualities of a robot. It is possible a number of other perceptions changed, but my instruments were not broad enough to cover those areas. I found that gathering rich qualitative feedback coupled with qualitative analysis was very useful for exploring a user's experience but developing more targeted tools and measurement approaches can be useful to the field as a whole.

Video games themselves can be a source of such measurements, or at least provide a base to extend to teleoperation. Hochleitner *et al.* provide a good recent overview of evaluating user experience in video games (Hochleitner et al. 2015). Their list of heuristics is long and broad, and they note are even sometimes specific to certain video game genres, further demonstrating the difficulty of evaluating user experience as measurements may be specific to the intended experience. As this thesis has demonstrated the applicability of techniques from video games to benefit teleoperation, future work may wish into teleoperation experience measurements may look to how video game user experience metrics themselves can be leveraged to better evaluate a teleoperation interaction design.

Another challenge I encountered at all stages of this thesis was creating metrics to evaluate teleoperator behaviour, and I believe this needs to be considered in all future teleoperation work continuing this line of research. Video games are often built to encourage certain styles of play, such as aggressive play, or strategic, etc. I similarly tried to encourage and measure an operator's tendency to drive safely or recklessly. In retrospect, however, my measurements of collisions and collisions over time, while useful in that it helped me understand if operators made fewer errors, did not provide enough power to understand how their behaviour changed. For example, a low number of collisions does not help me differentiate a confident and competent driver compared to a nervous and cautious driver. Future work may do trajectory analysis on the path by the operator or analysis of joystick input (e.g., frantic movements or smooth controls) may provide insight. Combining these with other descriptive measurements such as robot velocity, user measures such as stress indicators, or learning from other fields such as how driving motor vehicles is evaluated may provide more insight into how operators experience teleoperation and how that translates to changes in behaviour.

Experimental Design Considerations

Many choices go into designing an evaluation, and one common element in all of ours was to perform within-participant evaluation, where each operator used all of my teleoperation interface variations. This enabled participants to directly compare each in my open-ended feedback questionnaires, and further provided statistical power to factor out personal skill differences in my quantitative analyses. However, I question if participants had specific experiences they described in my data *because* the experiments were within-participant. In other words, would I be able to replicate my results in a between-participants design, and which should be recommended for future work?

Looking to the background research that informed my work, techniques such as priming (Yi 1990; Dijksterhuis and Bargh 2001) are impacted by prior experience, and user experience in general is understood to be temporal and contextual – when and in what context something is experienced is key to user experience (Forlizzi and Battarbee 2004; Hassenzahl and Tractinsky 2006). Therefore, I would expect my effects to change in between-participant evaluations, leading me to questions about the external validity of the effects I observed – as my work heavily relied on user experience, participants' responses are likely affected by their experiences with the previous conditions in the experiment. Due to the context-sensitivity of some user experiences, however, I hesitate to guess how experimental design changes would manifest in my results.

8.2 Looking Forward

In addition to the suggestion of continuing to explore and design more vide game-inspired interactions, there are two other important and broader areas of future work emerging from my research: user modelling and the ethics of experience manipulation.

8.2.1 Further Exploring Video Game-Inspired Teleoperation

One purpose of my framework of video game interaction (Chapter Six) is to act as a guide for potential future interaction designs that could be investigated for application in teleoperation. By abstracting existing video game designs, we can see a number of broad interaction styles that have yet to be researched in detail, such as different types of messages like interactable messages (for example, communicating what a robot could interact with in the environment and how the operator could perform that interaction). Identifying these less explored ideas, examining existing implementations of the designs in video games, and then adapting to teleoperation with iterative user evaluation, will likely yield many more ways to improve teleoperation performance and experience.

In particular, continuing investigation of video game interfaces that leverage different aspects of human psychology is a promising future work direction. One common example I witnessed was the manipulation of how resources were displayed to create specific responses in players, such as how health display may be non-linear: a health bar may show a small sliver of health left when in reality the player may still have much more left, creating the feeling of being in danger and promoting careful play. My prototypes throughout this thesis used different aspects of human psychology to modify different aspects of user experience, and I recommend pursuing this direction when analyzing and designing new video game-inspired teleoperation designs.

8.2.2 User-Modelling for Robot Awareness of Operator Experience

Most of my work and the related work focuses on the operator's awareness of the remote state: what may be interesting in the robot's environment, the robot's capabilities, and the state of the robot being driven. Situation awareness, however, also includes system awareness of the user (Endsley 2016). I believe that an important area of future work in teleoperation will be considering this information flow about the operator to the robot and system.

User modelling is the idea that I can build a model of a user (their physical, cognitive, emotional states, etc.), and use that as input into an algorithm to adjust the system. One could imagine a sensors that read social signals such as an operator's body language, or emotional state with galvanic skin response measurements, to better understand an operator's state and make adjustments accordingly. For example, a system may detect that a user is becoming tired and adjust or reduce extra interface components that may be creating too high of a cognitive load, or add interfaces that raise an operator's confidence when they are becoming frustrated.

I found estimating or measuring the user's state to steer them towards a desired experience is also an approach leveraged in video games. Games have systems that increase the amount of automated assistance that is offered [e.g. auto-aim, Resident Evil 4, Capcom, 2005], intelligently pace rest areas and encounter difficulties [Left4Dead 2, Valve, 2009], or offer layered complexity that take into account player improvement as the game progresses [tutorials in Kingdom Hearts Series, Square Enix, 2002-2019, or increased equipment choices as the game progresses in The Legend of Zelda Series, Nintendo, 1986-2018]. Similar teleoperation systems may be implemented to dynamically increase robot control assistance from system algorithms or help build better training programs to master teleoperation systems when the user is detected to be frustrated. To my knowledge, ideas like user-modelling that take the user state into account are an under-served research area in teleoperation and is an interesting direction for future work.

8.2.3 The Ethics of Manipulating User Experience

Taking a different lens to my research, it may appear that my work is manipulative: I used techniques similar to misdirection in magic to draw operator attention to specific events, deceived operators outright to build specific assumptions of robot abilities, and used ingrained social communication strategies to make operators feel worse about their teleoperation performance. Should we encourage research in this direction? I found only small discussion, and a lack of tools such as heuristics or frameworks, describing and exploring the ethics of manipulative robotic interfaces. I believe such discourse, guidelines, and ethical evaluation toolkits are an important topic for future work.

Looking to my inspiration, video games, I notice a distinct difference with teleoperation when considering the application of these techniques. Chiefly, video games are known to be virtual, and are generally viewed as entertainment, and people may go in with the intent to suspend their disbelief and *intend* to be deceived by the game for amusement purposes. People may not have that intent when operating a robot, and further may not even be aware an attempt to manipulate their experience is occurring.

While I could imagine systems that calm down operators in tense search-and-rescue scenarios, I can also imagine uses of priming to convince a potential customer that a robot performs better than it does in real life, potentially manipulating them to make an expensive purchase that does not fulfill their needs. Like much of technology, the design techniques I explored here can be used for a variety of purposes; the ethics of such potentially manipulative technology is an ongoing discussion in the social human-robot interaction community (Sanoubari et al. 2019; Winkle 2019; Henkel and Bethel 2017), and I show that teleoperation designers also need to be aware of these potential scenarios as well. Indeed, I think my research improves awareness that designs like those in this thesis can be used to encourage people to act or feel in certain ways enables designers to control these effects and not let such manipulation happen by accident.

8.3 Future Work Summary

Moving forward, my research can be extended by exploring more video game interaction, including those in genres radically different than teleoperation, and leads me to further interesting research topics in the use of user modelling and ethics in teleoperation. For this research, I recommend pursuing new methods and measurement tools to define and quantify both user experience and behaviour in teleoperation. My research points to a new path forward in teleoperation, focusing on more user-centered measurements in addition to task performance.

CONTRIBUTIONS AND FINAL WORDS

Teleoperation has many potential applications in both the home, industry, and other areas such as search and rescue. However, safe and efficient teleoperation is not trivial for operators, and improved interaction design is one way to improve the difficulties operators encounter. In this thesis, I investigated how the human-robot interaction community can improve teleoperation interfaces by taking inspiration from video game interfaces, taking four different angles: directing operator attention, priming perceptions of robot capability, using social agents to influence operator experience, and surveying and classifying a breadth of video game interaction techniques into a framework. These four angles resulted in a number of contributions to teleoperation interaction design:

i) A set of novel concrete interaction techniques for teleoperation, including a range of designs, prototypes, and formal evaluations. This includes a set of attention-drawing designs, two methods for priming perceptions of operator safety, one method for shaping an operator's emotions during teleoperation, and the knowledge that reducing robot speed can improve operator safety with a nominal increase in task completion time.

ii) A framework that abstracts video game design in to more fundamental components. This enables the field to establish the similarity between the problem spaces and design goals in teleoperation and video game interaction design. It further provides ways to discuss the wide variety of game interaction designs and generate new interaction designs for teleoperation.

iii) The approach of using priming to influence teleoperator behavior and perception. This includes verbally and visually describing the robot's abilities, suggesting robot abilities through a tangible feel, or modifying the robot's driving profile to appear safer and easier to drive.

iv) The pioneering of how a social agent, inspired from video games, can be leveraged to improve teleoperation experience. This opens teleoperation interface design to take advantage of techniques from social human-robot interaction in addition to the social techniques in video games.

v) A set of design parameters for directing visual attention that can be used in other visual interface designs in teleoperation.

vi) A reflection on the importance of user experience in teleoperation, how experience can improve operator performance, and how experience can be shaped by video game interaction designs.

Much of my contribution is in my exploration of different high-level video game interaction design approaches and how I adapted them to teleoperation interfaces and experimentally evaluated their effects. In some cases, I contributed whole new approaches to teleoperation interface design, such as through influencing operator perceptions with priming, and integrating social techniques in video games and social human-robot interaction into teleoperation. These works demonstrate that video game interaction designs can improve the operator experience and performance by leveraging knowledge of human psychology for a range of effects, applications, and design approaches.

My game survey and analysis provide the human-robot interaction community with a bridge from the video game industry to teleoperation. In it, I developed tools and vocabulary to help researchers discuss and apply different video game interface designs to teleoperation by looking at the interaction goal and purpose on an abstract level. In addition, my abstraction enabled me to see how common design goals and information used in game interactions are shared with teleoperation, suggesting a broader applicability of my approach.

Taking into these contributions into account, I can comment on the questions posed in my research objectives:

1) What types of interfaces exist in video games and how would they be beneficial to teleoperation?

   Almost every interface I observed had the basis for being applied beneficially to teleoperation (contributions i-vi). Even less obviously applicable designs, like the social-oriented interfaces investigated in Chapter Five were applicable once I understood the psychological and sociological principles behind the original design (contribution iv). Looking across dozens of video games belonging to first- or third-person action, role-playing, shooting, and strategy games, I can conclude that video game interfaces from these genres have wide-reaching potential and applicability to teleoperation problems (contributions i, ii).

2) How do human factors or psychological mechanisms used in video games affect teleoperation?

   I found video games leverage a variety of psychological and sociological phenomena in their designs (contribution ii). These include physiologically-driven actions (contribution v), cognitively driven perceptions (contribution iii), and automatic social responses (contribution iv), which may affect operator performance or experience. Thus, video game interfaces leverage a broad range of human mechanisms, and, when applied to teleoperation, can affect key performance metrics like safety.

3) How do game user experience goals change the teleoperation user experience?

   Research has noted, and I further observed, that games broadly strive to create certain user experiences with their interfaces (contribution ii, vi). I also noted experiential changes in each of my targeted explorations, including feeling tired or distracted (Chapter Three), perceiving a robot as faster or easier to control (Chapter Four), or feeling different emotions after using a robot (Chapter Five). Thus, video game interfaces can affect the teleoperation user experience in a variety of ways, and these experience changes may even drive operator performance changes as well (contribution vi).

I can conclude that video game interaction design shares a number of problems and goals with teleoperation, and that the interactions designed to improve video games can also improve the user experience and operator driving behaviour in teleoperation. Throughout exploring this idea, I provided a number of experimentally verified novel interaction techniques and directions for future exploration. Further, I found that designing to improve user experience, like video games frequently do, can similarly improve the teleoperation experience and can even improve operator task performance. By considering my own video game-inspired designs and my survey together, this thesis provides the first thorough and multifaceted analysis of the overarching strategy of learning from video games for teleoperation interface techniques.

REFERENCES

Aarts, Henk, and Ap Dijksterhuis. 2003. "The Silence of the Library: Environment, Situational Norm, and Social Behavior." *Journal of Personality and Social Psychology* 84 (1): 18–28.

Abrams, Richard A., and Shawn E. Christ. 2003. "Motion Onset Captures Attention." *Psychological Science* 14 (5): 427–432.

Abrams, Richard A., and Shawn E. Christ. 2006. "Motion Onset Captures Attention : A Rejoinder to Franconeri and Simons (2005)." *Perception & Psychophysics* 63130 (1): 114–117.

Ackerman, Joshua M, Christopher C Nocera, and John a Bargh. 2010. "Incidental Haptic Sensations Influence Social Judgments and Decisions." *Science* 328 (5986): 1712–1715.

Admoni, Henny, and Brian Scassellati. 2017. "Social Eye Gaze in Human-Robot Interaction: A Review." *Journal of Human-Robot Interaction* 6 (1): 25.

Andersen, Peter Bogh. 2001. "What Semiotics Can and Cannot Do for HCI." *Knowledge-Based Systems* 14 (8): 419–424.

Antin, Judd, and Ef Churchill. 2011. "Badges in Social Media: A Social Psychological Perspective." In *Human Factors in Computing Systems,* 1–4. ACM.

Atkinson, David J., and Micah H. Clark. 2014. "Methodology for Study of Human-Robot Social Interaction in Dangerous Situations." In *Human-Agent Interaction,* 371–376. New York, New York, USA, New York, USA: ACM Press.

Atrey, Pradeep K., M. Anwar Hossain, and Abdulmotaleb El Saddik. 2008. "Automatic Scheduling of CCTV Camera Views Using a Human-Centric Approach." In *IEEE International Conference on Multimedia and Expo*, 325–328. IEEE.

Babu, James. 2012. "Video Game HUDs: Information Presentation and Spatial Immersion James."

Bailey, Brian P., and Shamsi T. Iqbal. 2008. "Understanding Changes in Mental Workload during Execution of Goal-Directed Tasks and Its Application for Interruption Management." *ACM Transactions on Computer-Human Interaction* 14 (4): 1–28.

Balaguer, Carlos, Antonio Giménez, Alberto Jardón, Raúl Correal, Santiago Martínez, Angelo M. Sabatini, and Vincenzo Genovese. 2007. "Proprio and Teleoperation of a Robotic System for Disabled Persons' Assistance in Domestic Environments." *Springer Tracts in Advanced Robotics* 31: 415–427.

Banakou, Domna, Raphaela Groten, and Mel Slater. 2013. "Illusory Ownership of a Virtual Child Body Causes Overestimation of Object Sizes and Implicit Attitude Changes." *Proceedings of the National Academy of Sciences of the United States of America* 110 (31): 12846–12851.

Banh, Amy, Daniel J D.J. Rea, James E J.E. Young, and Ehud Sharlin. 2015. "Inspector Baxter : The Social Aspects of Integrating a Robot as a Quality Inspector in an Assembly Line." In *Human-Agent Interaction*. ACM.

Barber, Christine M., Robin J. Shucksmith, Bruce MacDonald, and Burkhard C. Wünsche. 2010. "Sketch-Based Robot Programming." *International Conference Image and Vision Computing New Zealand*.

Bargh, John A., Mark Chen, and Lara Burrows. 1996. "Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype Activation on Action." *Journal of Personality and Social Psychology* 71 (2): 230–244.

Barsade, Sigal G. 2002. "The Ripple Effect: Emotional Contagion and Its Influence on Group Behavior." *Administrative Science Quarterly* 47 (4): 644.

Bartneck, Christoph, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. "Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots." *International Journal of Social Robotics*.

Bartneck, Christoph, Michel van der Hoek, Omar Mubin, and Abdullah Al Mahmud. 2007. "'Daisy, Daisy, Give Me Your Answer Do!'" *Proceeding of the ACM/IEEE International Conference on Human-Robot Interaction - HRI '07*, no. 2007: 217.

Basu, Chandrayee, Qian Yang, David Hungerman, Mukesh Singhal, and Anca D. Dragan. 2017. "Do You Want Your Autonomous Car To Drive Like You?" In *Human-Robot Interaction*, 417–425. New York, New York, USA: ACM/IEEE.

Becker, Suzanna, Morris Moscovitch, Marlene Behrmann, and Steve Joordens. 1997. "Long-Term Semantic Priming: A Computational Account and Empirical Evidence." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23 (5): 1059–1082.

Berg, Bruce L. 1989. "Open Coding." In *Qualitative Research Methods for the Social Sciences*, 364–366.

Beyer, Hugh, and Karen Holtzblatt. 1997. *Contextual Design: Defining Customer-Centered Systems*. Elsevier.

Blanco, Myra, Wayne J. Biever, John P. Gallagher, and Thomas A. Dingus. 2006. "The Impact of Secondary Task Cognitive Processing Demand on Driving Performance." *Accident Analysis & Prevention* 38 (5): 895–906.

Blascovich, J. J., and E. S. Katkin. 1993. *Cardiovascular Reactivity to Psychological Stress & Disease.* Edited by James J. Blascovich and Edward S. Katkin. Washington: American Psychological Association.

Blommer, Mike, Reates Curry, Radhakrishnan Swaminathan, Louis Tijerina, Walter Talamonti, and Dev Kochhar. 2017. "Driver Brake vs. Steer Response to Sudden Forward Collision Scenario in Manual and Automated Driving Modes." *Transportation Research Part F: Traffic Psychology and Behaviour* 45. Elsevier Ltd: 93–101.

Borenstein, Jason, and Yvette Pearson. 2011. "16 Robot Caregivers: Ethical Issues across the Human Lifespan." In *Robot Ethics: The Ethical and Social Implications of Robotics*, 251. MIT Press.

Borji, Ali, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. 2015. "Salient Object Detection: A Benchmark." *IEEE Transactions on Image Processing* 24 (12): 5706–5722.

Bowman, D.A., David Koller, and L.F. Hodges. 1997. "Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques." In *Proceedings of IEEE 1997 Annual International Symposium on Virtual Reality*, 45-52,. IEEE Comput. Soc. Press.

Bradley, Adam, Cayley MacArthur, Mark Hancock, and Sheelagh Carpendale. 2015. "Gendered or Neutral ? Considering the Language of HCI." In *Graphics Interfaces*. NS: ACM Press.

Brandstetter, Jürgen, Clayton Beckner, Eduardo Benítez Sandoval, and Christoph Bartneck. 2017. "Persistent Lexical Entrainment in HRI." In *Human-Robot Interaction*, 63–72. ACM.

Breazeal, Cynthia. 2003. "Toward Sociable Robots." In *Robotics and Autonomous Systems*, 42:167–175.

Breazeal, Cynthia, Cory D Kidd, Andrea Lockerd Thomaz, Guy Hoffman, and Matt Berlin. 2005. "Effects of Nonverbal Communication on Efficiency and Robustness of Human-Robot Teamwork.Pdf." *IEEE International Conference on Intelligent Robots & Systems*.

Bridgeman, B., D. Hendry, and L. Stark. 1975. "Failure to Detect Displacement of the Visual World during Saccadic Eye Movements." *Vision Research* 15 (6): 719–722.

Briggs, Gordon, and Matthias Scheutz. 2014. "How Robots Can Affect Human Behavior: Investigating the Effects of Robotic Displays of Protest and Distress." *International Journal of Social Robotics* 6 (3): 343–355.

Bruce, Neil D B, Christopher Catton, and Sasa Janjic. 2016. "A Deeper Look at Saliency: Feature Contrast, Semantics, and Beyond." In *The IEEE Conference on Computer Vision and Pattern*

Recognition (CVPR).

Burgess, Joel (Bethesda). 2014. "Level Design in a Day: How We Used Iterative Level Design to Ship Skyrim and Fallout 3." In *Game Developers Conference*.

Burke, Moira, Anthony Hornof, Erik Nilsen, and Nicholas Gorman. 2005. "High-Cost Banner Blindness : Ads Increase Perceived Workload , Hinder Visual Search , and Are Forgotten." *ACM Transactions on Computer-Human Interaction* 12 (4): 423–445.

Buxton, William. 1988. "Taxonomies of Input." *Haptic Input*, 1–16.

Cai Meng, Tianmiao Wang, Wusheng Chou, Sheng Luan, Yuru Zhang, and Zengmin Tian. 2004. "Remote Surgery Case: Robot-Assisted Teleneurosurgery." In *Robotics and Automation*, 819-823 Vol.1. IEEE.

Candland, David (Bungie). 2016. "Tenacious Design and The Interface of 'Destiny.'" In *Game Developers Conference*.

Caroux, Loïc, and Katherine Isbister. 2016. "Influence of Head-up Displays' Characteristics on User Experience in Video Games." *International Journal of Human Computer Studies* 87. Elsevier: 65–79.

Carrol, John Ni, and John C Thoma. 1988. "Fun" 19 (3): 21–24.

Chalfoun, Pierre, and Claude Frasson. 2011. "Subliminal Cues While Teaching: HCI Technique for Enhanced Learning." *Advances in Human-Computer Interaction* 2011: 1–15.

Cheesman, Jim, and Philip M. Merikle. 1984. "Priming with and without Awareness." *Perception & Psychophysics* 36 (4): 387–395.

Chen, Jessie Y. C., Ellen C. Haas, and Michael J. Barnes. 2007. "Human Performance Issues and User Interface Design for Teleoperated Robots." *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)* 37 (6): 1231–1245.

Chen, Jessie Y C, Michael J. Barnes, and Michelle Harper-Sciarini. 2011. "Supervisory Control of Multiple Robots: Human-Performance Issues and User-Interface Design." *IEEE Transactions on Systems, Man and Cybernetics Part C: Applications and Reviews* 41 (4): 435–454.

Chun, Marvin M. 2000. "Contextual Cueing of Visual Attention." *Trends in Cognitive Sciences* 4 (5): 170–178.

"Clearpath Robotics: Controllers." 2017.

Coeckelbergh, Mark. 2011. "Talking to Robots: On the Linguistic Construction of Personal Human-Robot Relations." *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST* 59 LNICST: 126–129.

Corbin, Juliet, and Anselm Strauss. 1990. "Grounded Theory Research: Procedures, Canona and Evaluative Criteria." *Zeitschrift Fur Sociologie* 19 (6): 418–427.

Cormier, D., G. Newman, M. Nakane, and J. E. Young. 2013. "Would You Do as a Robot Commands? An Obedience Study for Human-Robot Interaction." In *In Proc. of the First International Conference on Human-Agent Interaction, IHAI'13*.

Cotter, TS. 2014. "Structural Analysis of Hfacs in Unmanned and Manned Air Vehicles." *Of the International Annual Conference of the …*.

Csikszentmihalyi, M. 1990. *Flow. The Psychology of the Optimal Experience.* Harper & Row.

Cui, Rongxin, Bo Gao, and Ji Guo. 2012. "Pareto-Optimal Coordination of Multiple Robots with Safety Guarantees." *Autonomous Robots* 32 (3): 189–205.

de Souza, C.S. 1993. "The Semiotic Engineering of User Interface Language Design." *Man-Machine Studies* 39: 753–773.

Deterding, Sebastian. 2012. "Gamification: Designing for Motivation." *Interactions* 19 (4): 14.

Deterding, Sebastian, Dan Dixon, Rilla Khaled, and Lennart Nacke. 2011. "From Game Design Elements to Gamefulness." In *Envisioning Future Media Environments - MindTrek '11*, 9–

11.

Dijksterhuis, a, and Ja Bargh. 2001. “The Perception-Behavior Expressway: Automatic Effects of Social Perception on Social Behavior.” *Advances in Experimental Social Psychology* 33: 1–38.

Dix, Alan, Janet E. Finlay, Gregory D. Abowd, and Russell Beale. 2003. *Human-Computer Interaction*. 3rd ed. Prentice Hall PTR.

Doherty, R. 1997. “The Emotional Contagion Scale: A Measure of Individual Differences.” *Journal of Nonverbal Behavior* 21 (2): 131–154.

Donald, Fiona, Craig Donald, and Andrew Thatcher. 2015. “Work Exposure and Vigilance Decrements in Closed Circuit Television Surveillance.” *Applied Ergonomics* 47 (January 2016): 220–228.

Dourish, Paul. 2001. “Seeking a Foundation for Context-Aware Computing Corresponding Author ’ s Contact Information : Department of Information.” *Human-Computer Interaction* 16 (2): 229–241.

Doyen, Stéphane, Olivier Klein, Cora Lise Pichon, and Axel Cleeremans. 2012. “Behavioral Priming: It’s All in the Mind, but Whose Mind?” *PLoS ONE* 7 (1).

Drury, J.L., J. Scholtz, and H.a. Yanco. 2003. “Awareness in Human-Robot Interactions.” *IEEE International Conference on Systems, Man and Cybernetics.* 1 (October).

Edsinger, Aaron, and Charles C. Kemp. 2007. “Human-Robot Interaction for Cooperative Manipulation: Handing Objects to One Another.” In *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, 1167–1172.

Elvik, Rune. 2007. “Effects on Road Safety of Converting Intersections to Roundabouts: Review of Evidence from Non-U.S. Studies.” *Transportation Research Record: Journal of the Transportation Research Board* 1847 (1): 1–10.

Endsley, Mica R. 1988. "Design and Evaluation for Situation Awareness Enhancement." *Proceedings of the Human Factors Society Annual Meeting* 32 (2): 97–101.

Endsley, Mica R. 2016. *Designing for Situation Awareness: An Approach to User-Centered Design*. Second. CRC Press.

Erden, Mustafa Suphi. 2013. "Emotional Postures for the Humanoid-Robot Nao." *International Journal of Social Robotics* 5 (4): 441–456.

Ermi, Laura, and Frans Mäyrä. 2005. "Fundamental Components of the Gameplay Experience: Analysing Immersion." *Changing Views: Worlds in Play*, 15–27.

Essential Facts about the Computer and Video Game Industry. 2015.

Eyssartier, Chloé, Sébastien Meineri, and Nicolas Gueguen. 2017. "Motorcyclists' Intention to Exceed the Speed Limit on a 90km/h Road: Effect of the Type of Motorcycles." *Transportation Research Part F: Traffic Psychology and Behaviour* 45: 183–193.

Feldmaier, Johannes, Martin Stimpfl, and Klaus Diepold. 2017. "Development of an Emotion-Competent SLAM Agent." In *Human-Robot Interaction*, 1–9. ACM Press.

Forlizzi, Jodi. 2007. "How Robotic Products Become Social Products: An Ethnographic Study of Cleaning in the Home." *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 129–136.

Forlizzi, Jodi, and Katja Battarbee. 2004. "Understanding Experience in Interactive Systems." *Proceedings of the 2004 Conference on Designing Interactive Systems Processes, Practices, Methods, and Techniques - DIS '04*, 261.

Forlizzi, Jodi, John Zimmerman, Vince Mancuso, and Sonya Kwak. 2007. "How Interface Agents Affect Interaction between Humans and Computers." *Proceedings of the 2007 Conference on Designing Pleasurable Products and Interfaces - DPPI '07*, 209.

Forster, Kenneth I, and Chris Davis. 1984. "Repetition Priming and Frequency Attenuation in

Lexical Access" 10 (4): 680–698.

Franconeri, Steven L, and Daniel J Simons. 2005. "The Dynamic Events That Capture Visual Attention: A Reply to Abrams and Christ (2005)." *Perception & Psychophysics* 67 (6): 962–966.

Fredrickson, B. 1998. "What Good Are Positive Emotions?" *Review of General Psychology* 2 (3): 300–319.

Fredrickson, Barbara L. 2001. "The Role of Positive Emotions in Positive Psychology: The Broaden-and-Build Theory of Positive Emotions Perspectives on Emotions and Affect." *American Psychologist* 56 (3): 218–226.

Fukatsu, Shinji, Yoshifumi Kitamura, Toshihiro Masaki, and Fumio Kishino. 1998. "Intuitive Control of 'Bird's Eye' Overview Images for Navigation in an Enormous Virtual Environment." In *ACM Virtual Reality Software and Technology*, 67–76.

Fuller, Ray. 2005. "Towards a General Theory of Driver Behaviour." *Accident Analysis and Prevention* 37 (3): 461–472.

Gaver, William W. 1991. "Technology Affordances." In *Human Factors in Computing Systems*, 79–84. ACM.

Geiskkovitch, Denise Y, Derek Cormier, Stela H Seo, and James E Young. 2016. "Please Continue, We Need More Data: An Exploration of Obedience to Robots." *Journal of Human-Robot Interaction* 5 (1): 82–99.

Gibson, James J. 2015. *The Ecological Approach to Visual Perception: Classic Edition*. Classic. Psychology Press.

Giese, Stefanie, David Carr, and Javaan Chahl. 2013. "Implications for Unmanned Systems Research of Military UAV Mishap Statistics." *IEEE Intelligent Vehicles Symposium, Proceedings*, no. Iv: 1191–1196.

Glas, Dylan F., Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2012. "Teleoperation of Multiple Social Robots." *IEEE Transactions on Systems, Man, and Cybernetics* 42 (3): 530–544.

Glaser, B. G., and A. L. Strauss. 2017. *Discovery of Grounded Theory: Strategies for Qualitative Research*. Routledge.

Gleeson, Brian, Karon Maclean, Amir Haddadi, Elizabeth Croft, and Javier Alcazar. 2013. "Gestures for Industry: Intuitive Human-Robot Communication from Human Observation." In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI '13)*, 349–356. Piscataway, NJ, USA: IEEE Press.

Goodrich, M. a., J. W. Crandall, and E. Barakova. 2013. "Teleoperation and Beyond for Assistive Humanoid Robots." *Reviews of Human Factors and Ergonomics* 9 (1): 175–226.

Goodrich, M.A., and D.R. Olsen. 2003. "Seven Principles of Efficient Human Robot Interaction." In *Systems, Man and Cybernetics*, 4:3942–3948. IEEE.

Goodrich, Michael A, and Alan C Schultz. 2008. *Human–Robot Interaction: A Survey*. *Foundations and Trends® in Human–Computer Interaction*. Vol. 1.

Gopinath, Deepak, Siddarth Jain, and Brenna D. Argall. 2017. "Human-in-the-Loop Optimization of Shared Autonomy in Assistive Robotics." *IEEE Robotics and Automation Letters* 2 (1): 247–254.

Grandjean, E. 1980. *Fitting the Task to the Man: An Ergonomic Approach*. Taylor & Francis.

Groeger, J.A., and J.A. Rothengatter. 1998. "Traffic Psychology and Behaviour." *Transportation Research Part F: Traffic Psychology and Behaviour* 1 (1): 1–9.

Groeger, John A. 2002. "Trafficking in Cognition: Applying Cognitive Psychology to Driving." *Transportation Research Part F: Traffic Psychology and Behaviour* 5 (4): 235–248.

Gross, James J. 2002. "Emotion Regulation : Affective , Cognitive , and Social Consequences," 281–291.

Grossman, Tovi, George Fitzmaurice, and Ramtin Attar. 2009. "A Survey of Software Learnability: Metrics, Methodologies and Guidelines." *Proceedings of the 27th International Conference on Human Factors in Computing Systems (CHI'09)*, 649–658.

Grundberg, Sven, and Jens Hansegard. 2014. "Women Now Make Up Almost Half of Gamers." *Wall Street Journal*, August 20.

Guo, Cheng, James Everett Young, and Ehud Sharlin. 2009. "Touch and Toys." In *Human Factors in Computing Systems*, 491. New York, New York, USA: ACM Press.

Guo, Hongling, Heng Li, Greg Chan, and Martin Skitmore. 2012. "Using Game Technologies to Improve the Safety of Construction Plant Operations." *Accident Analysis and Prevention* 48. Elsevier Ltd: 204–213.

Gustafson, Sean, Patrick Baudisch, Carl Gutwin, and Pourang Irani. 2008. "Wedge: Clutter-Free Visualization of off-Screen Locations." In *Human Factors in Computing Systems (CHI)*, 787. New York, New York, USA: ACM Press.

Hacinecipoglu, Akif, E. Ilhan Konukseven, and A. Bugra Koku. 2013. "Evaluation of Haptic Feedback Cues on Vehicle Teleoperation Performance in an Obstacle Avoidance Scenario." In *2013 World Haptics Conference, WHC 2013*.

Hamari, Juho, Jonna Koivisto, and Harri Sarsa. 2014. "Does Gamification Work? - A Literature Review of Empirical Studies on Gamification." *Proceedings of the Annual Hawaii International Conference on System Sciences*, no. JANUARY: 3025–3034.

Hancock, P. a., D. R. Billings, K. E. Schaefer, J. Y. C. Chen, E. J. de Visser, and R. Parasuraman. 2011. "A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction." *Human Factors: The Journal of the Human Factors and Ergonomics Society* 53 (5): 517–527.

Harrison, Lane, Drew Skau, Steven Franconeri, Aidong Lu, and Remco Chang. 2013. "Influencing Visual Judgment through Affective Priming," 2949.

Hart, Sandra, G. 2006. "NASA-Task Load Index (NASA-TLX); 20 Years Later." *Human Factors and Ergonomics Society Annual Meting*, 904–908.

Hart, Sandra G., and Lowell E. Staveland. 1988. "Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research." In *Human Mental Workload*, 139–183.

Hashimoto, Sunao, Akihiko Ishida, Masahiko Inami, and Takeo Igarash. 2011. "TouchMe: An Augmented Reality Based Remote Robot Manipulation." *International Conference on Artificial Reality and Telexistence*, 1–6.

Hassenzahl, Marc, and Noam Tractinsky. 2006. "User Experience - A Research Agenda." *Behaviour and Information Technology* 25 (2): 91–97.

Hatfield, Elaine, John Cacioppo, and Richard Rapson. 1994. "Emotional Contagion." *Current Directions in Psychological Science*.

Henkel, Zachary, and Cindy L. Bethel. 2017. "A Robot Forensic Interviewer." In *Human-Robot Interaction*, 10–20. New York, New York, USA: ACM Press.

Herr, Paul M. 2012. "Priming Price: Prior Knowledge and Context Effects." *Journal of Consumer Research* 16: 1–10.

Hochleitner, Christina, Wolfgang Hochleitner, Cornelia Graf, and Manfred Tscheligi. 2015. "A Heuristic Framework for Evaluating User Experience in Games." In *Game User Experience Evaluation*, 87–111.

Horvitz, Eric, Andy Jacobs, and David Hovel. 1999. "Attention-Sensitive Alerting." In *Uncertainty in Artificial Intelligence*, 305–313.

Horvitz, Eric, Carl Kadie, Tim Paek, and David Hovel. 2003. "Models of Attention in Computing and Communication." *Communications of the ACM* 46 (3): 52.

Howard, Christina J, Tomasz Troscianko, Iain D Gilchrist, Ardhendu Behera, and David C Hogg. 2009. "Searching for Threat : Factors Determining Performance during CCTV Monitoring."

Security, 1–7.

Iacoboni, Marco. 2009. “Imitation, Empathy, and Mirror Neurons.” *Annual Review of Psychology* 60 (1): 653–670.

Iacovides, Ioanna, Anna Cox, Richard Kennedy, Paul Cairns, and Charlene Jennett. 2015. “Removing the HUD.” *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play - CHI PLAY ’15*, 13–22.

“IRobot 510 PackBot.” 2015.

Isen, A. 1987. “Positive Affect Facilitates Creative Problem Solving.” *Personality and Social Psychology Bulletin1* 52 (6): 1122–1131.

Jennett, Charlene, Anna L. Cox, Paul Cairns, Samira Dhoparee, Andrew Epps, Tim Tijs, and Alison Walton. 2008. “Measuring and Defining the Experience of Immersion in Games.” *International Journal of Human Computer Studies* 66 (9): 641–661.

Johnson, Daniel, and Janet Wiles. 2003. “Effective Affective User Interface Design in Games.” *Ergonomics* 46 (13–14): 1332–1345.

Jonah, Brian A., Rachel Thiessen, and Elaine Au-Yeung. 2001. “Sensation Seeking, Risky Driving and Behavioral Adaptation.” *Accident Analysis and Prevention* 33 (5): 679–684.

Jung, Malte F., Nikolas Martelaro, and Pamela J. Hinds. 2015. “Using Robots to Moderate Team Conflict.” In *Human-Robot Interaction*, 229–236.

Kahn, Peter H. Jr., Takayuki Kanda, Hiroshi Ishiguro, Brian T. Gill, Jolina H. Ruckert, Solace Shen, Heather E. Gary, et al. 2012. “Do People Hold a Humanoid Robot Morally Accountable for the Harm It Causes?,” 33–40.

Kahn, Peter H., Takayuki Kanda, Hiroshi Ishiguro, Brian T. Gill, Solace Shen, Heather E. Gary, and Jolina H. Ruckert. 2015. “Will People Keep the Secret of a Humanoid Robot?” In *Human-Robot Interaction* -, 173–180.

Kahn, Peter H, Takayuki Kanda, Hiroshi Ishiguro, Nathan G Freier, Rachel L Severson, Brian T Gill, Jolina H Ruckert, and Solace Shen. 2012. "'Robovie, You'll Have to Go into the Closet Now': Children's Social and Moral Relationships with a Humanoid Robot." *Developmental Psychology* 48 (2): 303–314.

Kalkofen, Denis, Eduardo Veas, Stefanie Zollmann, Markus Steinberger, and Dieter Schmalstieg. 2013. "Adaptive Ghosted Views for Augmented Reality." *IEEE International Symposium on Mixed and Augmented Reality*, no. October: 1–9.

Kanade, T., J. F. Cohn, and Y Tian. 2000. "Comprehensive Database for Facial Expression Analysis." In *Automatic Face and Gesture Recognition*. IEEE.

Kato, Jun, and Takeo Igarashi. 2014. "VisionSketch: Integrated Support for Example-Centric Programming of Image Processing Applications." In *Graphics Interface*, 115–122. ACM.

Kato, Jun, Daisuke Sakamoto, and Takeo Igarashi. 2013. "Picode: Inline Photos Representing Posture Data in Source Code." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, 3097–3100. New York, New York, USA: ACM Press.

Kelley, Harold H. 1950. "The Warm-Cold Variable in First Impressions of Persons." *Journal of Personality* 18 (4): 431–439.

Keyes, Brenden, Robert Casey, Holly a Yanco, Bruce a Maxwell, and Yavor Georgiev. 2006. "Camera Placement and Multi-Camera Fusion for Remote Robot Operation." *Proceedings of the IEEE International Workshop on Safety, Security and Rescue Robotics*, 22–24.

Kirman, Ben, Conor Linehan, Shaun Lawson, and Dan O'Hara. 2013. "CHI and the Future Robot Enslavement of Humankind: A Retrospective." *CHI '13 Extended Abstracts on Human Factors in Computing Systems on - CHI EA '13*, 2199.

Klein, Raymond M, W Joseph Macinnes, Raymond M Klein, and W Joseph Macinnes. 1999. "Inhibition of Return Is a Foraging Facilitator," 346–352.

Klemmer, Scott R, Björn Hartmann, and Leila Takayama. 2006. "How Bodies Matter." In

Proceedings of the 6th ACM Conference on Designing Interactive Systems - DIS '06, 140. New York, New York, USA: ACM Press.

Koeffel, Christina, Wolfgang Hochleitner, Jakob Leitner, Michael Haller, Arjan Geven, and Manfred Tscheligi. 2010. “Using Heuristics to Evaluate the Overall User Experience of Video Games and Advanced Interaction Games.” In *Evaluating User Experience in Games*, 3–7.

Koffka, Kurt. 2013. *Principles of Gestalk Psychology*. Routledge.

Kolling, Andreas, Steven Nunnally, and Michael Lewis. 2012. “Towards Human Control of Robot Swarms.” In *Human-Robot Interaction*, 89–96. New York, New York, USA: ACM.

Krajník, Tomáš, Vojtěch Vonásek, Daniel Fišer, and Jan Faigl. 2011. “AR-Drone as a Platform for Robotic Research and Education.” In , 172–186.

Kristjánsson, Árni, Ómar I. Jóhannesson, and Ian M. Thornton. 2014. “Common Attentional Constraints in Visual Foraging.” Edited by Daniel Osorio. *PLoS ONE* 9 (6): e100752.

Kristoffersson, Annica, Silvia Coradeschi, and Amy Loutfi. 2013. “A Review of Mobile Robotic Telepresence.” *Advances in Human-Computer Interaction* 2013: 1–17.

Kubo, Norio, Volker Mellert, Reinhard Weber, and Jens Meschke. 2004. “Engine Sound Perception – Apart from so-Called Engine Order Analysis.” In *CFA/DAGA*, 867–868.

Kulik, a. 2009. “Building on Realism and Magic for Designing 3D Interaction Techniques.” *IEEE Computer Graphics and Applications* 29 (6): 22–33.

Labonte, Daniel, Patrick Boissy, and François Michaud. 2010. “Comparative Analysis of 3-D Robot Teleoperation Interfaces with Novice Users.” *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 40 (5): 1331–1342.

Langer, Rebecca, Mark Hancock, and Stacey D. Scott. 2015. “Suspenseful Design: Engaging Emotionally with Complex Applications through Compelling Narratives.” *Conference Proceedings - 2014 IEEE Games, Media, Entertainment Conference, IEEE GEM 2014*.

Lee, J. R., and C. I. Nass. 2010. "Trust in Computers: The Computers-Are-Social-Actors (CASA) Paradigm and Trustworthiness Perception in Human-Computer Communication." In *Trust and Technology in a Ubiquitous Modern Environment: Theoretical and Methodological Perspectives*, 1–15. IGI Global.

Leeper, a, Kaijen Hsiao, M Ciocarlie, L Takayama, and D Gossow. 2012. "Strategies for Human-in-the-Loop Robotic Grasping." *Human-Robot Interaction (HRI), 2012 7th ACM/IEEE International Conference On*, 1–8.

Leshed, Gilly, Theresa Velden, Oya Rieger, Blazej Kot, and Phoebe Sengers. 2008. "In-Car Gps Navigation: Engagement with and Disengagement from the Environment." *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, 1675–1684.

Levy-Tzedek, Shelly, Sigal Berman, Yehuda Stiefel, Ehud Sharlin, James Young, and Daniel Rea. 2017. "Robotic Mirror Game for Movement Rehabilitation." *International Conference on Virtual Rehabilitation, ICVR* 2017-June.

Li, Wei, Tovi Grossman, and George Fitzmaurice. 2012. "GamiCAD: A Gamified Tutorial System for First Time AutoCAD Users." *UIST*, 103–112.

Lim, Julian, Wen-chau Wu, Jiongjiong Wang, John A Detre, and David F Dinges. 2010. "Perfusion Study of the Time-On-Task Effect." *Neuroimage* 49 (4): 3426–3435.

Lindgaard, Gitte, Gary Fernandes, Cathy Dudek, and J. Brown. 2006. "Attention Web Designers: You Have 50 Milliseconds to Make a Good First Impression!" *Behaviour & Information Technology* 25 (2): 115–126.

Liu, Kexi, Daisuke Sakamoto, Masahiko Inami, Takeo Igarashi, and Baton Rouge. 2011. "Roboshop : Multi-Layered Sketching Interface for Robot Housework Assignment and Management." *Industrial Electronics*, 647–656.

Lucey, P., J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. 2010. "The Extended Cohn-Kanade Dataset (CK+): A Complete Expression Dataset for Action Unit and Emotion-Specified Expression." In *Workshop on CVPR for Human Communicative Behavior Analysis*.

Luria, Michal. 2018. “Designing Robot Behavior Based on Fictional Sidekick Characters.” In *Human-Robot Interaction*, 307–308. ACM.

Mack, Arien, and Irvine Rock. 1999. “Inattentional Blindness.” *Trends in Cognitive Sciences* 3 (1): 39.

MacLeod, Colin M. 1989. “Word Context during Initial Exposure Influences Degree of Priming in Word Fragment Completion.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15 (3): 398–406.

Maier, Jonathan R A, and Georges M. Fadel. 2009. “Affordance Based Design: A Relational Theory for Design.” *Research in Engineering Design* 20 (1): 13–27.

Malacria, Sylvain, Joey Scarr, Andy Cockburn, Carl Gutwin, and Tovi Grossman. 2013. “Skillometers: Reflective Widgets That Motivate and Help Users to Improve Performance.” *UIST*, 321–330.

Malone, Thomas W. 1981. “Toward a Theory of Intrinsically Motivating Instruction.” *Cognitive Science* 5 (4): 333–369.

Malone, Thomas W. 1982. “Heuristics for Designing Enjoyable User Interfaces: Lessons from Computer Games.” *Proceedings of the 1982 Conference on Human Factors in Computing Systems*, 63–68.

Manocha, K. A., N. Pernalete, and R. V. Dubey. 2001. “Variable Position Mapping Based Assistance in Teleoperation for Nuclear Cleanup.” *Robotics and Automation* 1: 374–379.

“Mario Kart 7 In-Game Statistics.” 2019.

Martin, P. Y. 1986. “Grounded Theory and Organizational Research.” *The Journal of Applied Behavioral Science* 22 (2): 141–157.

Mast, Marcus, Zdeněk Materna, Michal Španěl, Florian Weisshardt, Georg Arbeiter, Michael Burmester, Pavel Smrž, and Birgit Graf. 2015. “Semi-Autonomous Domestic Service Robots:

Evaluation of a User Interface for Remote Manipulation and Navigation With Focus on Effects of Stereoscopic Display." *International Journal of Social Robotics* 7 (2): 183–202.

Matthews, Tara Lynn. 2007. "Designing and Evaluating Glanceable Peripheral Displays." *Proceedings of the 6th Conference on Designing Interactive Systems*, 343–345.

Max M., North, North Sarah M., and Coble Joseph R. 1998. *Virtual Reality Therapy: An Effective Treatment for Phobias.* Amsterdam, Netherlands: IOS Press.

McCrickard, D. Scott, and C.M. Chewar. 2003. "Attuning Notification Design to User Goals and Attention Costs." *Communications of the ACM* 46 (3): 67–72.

Mcgrenere, Joanna, and Wayne Ho. 2000. "Affordances : Clarifying and Evolving a Concept." *Graphics Interface*, no. May: 1–8.

McIlroy, Rich C., Neville A. Stanton, and Louise Godwin. 2017. "Good Vibrations: Using a Haptic Accelerator Pedal to Encourage Eco-Driving." *Transportation Research Part F: Traffic Psychology and Behaviour* 46. Elsevier Ltd: 34–46.

Mekler, Elisa D, Julia Ayumi Bopp, Alexandre N Tuch, and Klaus Opwis. 2014. "A Systematic Review of Quantitative Studies on the Enjoyment of Digital Entertainment Games." *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems - CHI '14*, 927–936.

Michon, John A. 1985. "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?" *Human Behavior and Traffic Safety*, 485–520.

Miller, George. 1956. "The Magical Number Seven, Plus or Minus TwoSome Limits on Our Capacity for Processing Information." *Psychological Review* 101 (2): 343–352.

Mine, Mark, Arun Yoganandan, and Dane Coffey. 2014. "Making VR Work: Building a Real-World Immersive Modeling Application in the Virtual World." In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction - SUI '14*, 80–89. New York, New York, USA: ACM Press.

Mitra, Debanjan, and Peter N. Golder. 2006. “How Does Objective Quality Affect Perceived Quality? Short-Term Effects, Long-Term Effects, and Asymmetries.” *Marketing Science* 25 (3): 230–247.

Montagu, J. D., and E. M. Coles. 1966. “Mechanism and Measurement of the Galvanic Skin Response.” *Psychological Bulletin* 65 (5): 261–279.

Moody, Daniel L. 2009. “The ``Physics'' of Notations: Toward a Scientific Basis for ConstructingVisual Notations in Software Engineering.” *IEEE Transactions on Software Engineering* 35 (6): 756–779.

Moon, AJung, Chris A. C. Parker, Elizabeth A. Croft, and H. F. Machiel Van der Loos. 2013. “Design and Impact of Hesitation Gestures during Human-Robot Resource Conflicts.” *Journal of Human-Robot Interaction* 2 (3): 18–40.

Morris, JD. 1995. “Observations: SAM: The Self-Assessment Manikin; an Efficient Cross-Cultural Measurement of Emotional Response.” *Journal of Advertising Research*, no. December.

Murrell, KFH. 1965. *Ergonomics: Man in His Working Environment*. Chapman and Hall.

Nah, Fiona Fui Hoon, Brenda Eschenbrenner, Qing Zeng, Venkata Rajasekhar Telaprolu, and Sepandar Sepehr. 2014. “Flow in Gaming: Literature Synthesis and Framework Development.” *International Journal of Information Systems and Management* 1 (JANUARY 2014): 83–124.

Nasukawa, Tetsuya, and Jeonghee Yi. 2003. “Sentiment Analysis : Capturing Favorability Using Natural Language Processing.” *Proceedings of the 2nd International Conference on Knowledge Capture*, 70–77.

Neisser, U. 2014. *Cognitive Psychology: Classic Edition*. Psychology Press.

Nielsen, Curtis W., Michael A. Goodrich, and Robert W. Ricks. 2007. “Ecological Interfaces for Improving Mobile Robot Teleoperation.” *IEEE Transactions on Robotics* 23 (5): 927–941.

Norman, Donald A. 1983. "Design Rules Based on Analyses of Human Error." *Communications of the ACM* 26 (4): 254–258.

Norman, Donald A. 2004. *Emotional Design: Why We Love (or Hate) Everyday Things*. Civitas Books.

Norman, Donald A. 1986. "User-Centered Design: New Perspectives on Human-Computer Interaction." In *Cognitive Engineering*, edited by Norman, Donald A and Stephen Draper, 31–62. Erlbaum Associates.

Norman, Donald A. 2013. *The Design of Everyday Things*. *Human Factors and Ergonomics in Manufacturing*. Vol. 16.

Nunez, D, and E Blake. 2003. "Conceptual Priming as a Determinant of Presence in Virtual Environments." In *Computer Graphics, Virtual Reality and Visualisation in Africa*, 1:101–108. ACM.

Ogier, Harley, and Jim Buchan. 2017. "Exploring the Feasibility of Diegetic In-Game Store User Interfaces." In *Proceedings of the Australasian Computer Science Week Multiconference on - ACSW '17*, 1–10. New York, New York, USA: ACM Press.

Ohshima, Naoki, Keita Kimijima, Junji Yamato, and Naoki Mukawa. 2015. "A Conversational Robot with Vocal and Bodily Fillers for Recovering from Awkward Silence at Turn-Takings." In *International Workshop on Robot and Human Interactive Communication*, 2015-Novem:325–330. IEEE.

Olsen, Dan R., and Stephen Bart Wood. 2004. "Fan-out: Measuring Human Control of Multiple Robots." In *Human Factors in Computing Systems*, 6:231–238. New York, New York, USA: ACM Press.

Parasuraman, Raja, Thomas B. Sheridan, and Christopher D. Wickens. 2008. "Situation Awareness, Mental Workload, and Trust in Automation: Viable, Empirically Supported Cognitive Engineering Constructs." *Journal of Cognitive Engineering and Decision Making* 2 (2): 140–160.

Pascoe, Jason, Nick Ryan, and David Morse. 2000. “Using While Moving: HCI Issues in Fieldwork Environments.” *Transactions on Computer-Human Interaction (TOCHI) - Special Issue on Human-Computer Interaction with Mobile Systems* 7 (3): 417–437.

Paymans, Tim F., Jasper Lindenberg, and Mark Neerincx. 2004. “Usability Trade-Offs for Adaptive User Interfaces.” In *Proceedings of the 9th International Conference on Intelligent User Interface - IUI ’04*, 301. New York, New York, USA, New York, USA: ACM Press.

Peter, C., and R. Beale. 2008. “The Role of Affect and Emotion in HCI.” In *Affect and Emotion in Human-Computer Interaction - LNCS*. Springer.

Peter, Christian, and Antje Herbon. 2006. “Emotion Representation and Physiology Assignments in Digital Systems.” *Interacting with Computers* 18 (2): 139–170.

Picard, Rosalind W. 2000. *Affective Computing*. MIT Press.

Posner, Jonathan, James A Russell, Bradley S. Peterson, James A. Russellc, and Bradley S. Peterson. 2005. “The Circumplex Model of Affect: An Integrative Approach to Affective Neuroscience, Cognitive Development, and Psychopathology.” *Development and Psychopathology* 17 (3). Cambridge University Press: 715–734.

Posner, M. I., M. J. Nissen, and W. C. Ogden. 1978. “Attended and Unattended Processing Modes: The Role of Set for Spatial Location.” In *Modes of Perceiving and Processing Information*. Psychology Press.

Posner, Michael I. 1980. “Orienting of Attention.” *Q.J Exp.Psychol.* 32 (1): 3–25.

Postnikoff, Brittany, and Ian Goldberg. 2018. “Robot Social Engineering.” In *Companion of Human-Robot Interaction*, 313–314. New York, New York, USA: ACM Press.

Precht, Lisa, Andreas Keinath, and Josef F. Krems. 2017. “Effects of Driving Anger on Driver Behavior – Results from Naturalistic Driving Data.” *Transportation Research Part F: Traffic Psychology and Behaviour* 45. Elsevier Ltd: 75–92.

Preece, Jennifer, Helen Sharp, and Yvonne Rogers. 2015. *Interaction Design: Beyond Human-Computer Interaction, 4th Edition*. 4th ed. Wiley.

Pylyshyn, Z. W., and R. W. Storm. 1988. "Tracking Multiple Independent Targets: Evidence for a Parallel Tracking Mechanism." *Spatial Vision* 3 (3): 179–197.

Qin, Hua, Pei-Luen Patrick Rau, and Gavriel Salvendy. 2009. "Measuring Player Immersion in the Computer Game Narrative." *International Journal of Human-Computer Interaction* 25 (2): 107–133.

Quigley, M., M.A. Goodrich, and R.W. Beard. 2004. "Semi-Autonomous Human-UAV Interfaces for Fixed-Wing Mini-UAVs." *Intelligent Robots and Systems* 3: 2457–2462.

Radmard, Sina, Ajung Moon, and Elizabeth A Croft. 2015. "Interface Design and Usability Analysis for a Robotic Telepresence Platform." In *RO-MAN*, 6.

Rae, Irene, Leila Takayama, and Bilge Mutlu. 2013. "In-Body Experiences." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, 1921.

Razavi, Tiffani. 2001. "Self-Report Measures: An Overview of Concerns and Limitations of Questionnaire Use in Occupational Stress Research." University of Southampton.

Rea, D.J., T. Igarashi, and J.E. Young. 2014. "Paint Board- Prototyping Interactive Character Behaviors by Digitally Painting Storyboards." In *HAI 2014 - Proceedings of the 2nd International Conference on Human-Agent Interaction*, 315–322. ACM.

Rea, Daniel J., and James E. Young. 2019. "Methods and Effects of Priming a Teloperator's Perception of Robot Capabilities." In *Human-Robot Interaction,* 739–741. ACM/IEEE.

Rea, Daniel J, Mahdi Rahmani Hanzaki, Neil Bruce, and James E Young. 2017. "Tortoise and the Hare Robot Slow and Steady Almost Wins the Race , but Finishes More Safely." In *Robot and Human Interactive Communication (RO-MAN)*, 1–6. IEEE.

Rea, Daniel J, Stela H Seo, Neil Bruce, and James E Young. 2017. "Movers, Shakers, and Those

Who Stand Still: Visual Attention-Grabbing Techniques in Robot Teleoperation.” In *Human-Robot Interaction*, 398–407. New York, New York, USA: ACM/IEEE.

Rea, Daniel J, and James E Young. 2018. “It’s All in Your Head.” In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI ’18*, 32–40. New York, New York, USA, New York, USA: ACM Press.

Rea, Daniel J, and James E Young. 2019. “Backseat Teleoperator : Affective Feedback with On-Screen Agents to Influence Teleoperation.” *Human-Robot Interaction*, 19–28.

Recarte, Miguel a., and Luis M. Nunes. 1996. “Perception of Speed in an Automobile: Estimation and Production.” *Journal of Experimental Psychology: Applied* 2 (4): 291–304.

Reeves, B., and C. Nass. 1996. *How People Treat Computers, Television, and New Media like Real People and Places*.

Rensink, R. A., J. K. O’Regan, and J. J. Clark. 1996. “To See or Not to See: The Need for Attention to Perceive Changes in Scenes.” *Investigative Ophthalmology and Visual Science* 37 (3): 1–6.

Reveleau, Aurélien, François Ferland, Mathieu Labbé, Dominic Létourneau, and François Michaud. 2015. “Visual Representation of Sound Sources and Interaction Forces in a Teleoperation Interface for a Mobile Robot.” *Journal of Human-Robot Interaction* 4 (2): 1.

Richer, J., and J.L. Drury. 2006. “A Video Game-Based Framework for Analyzing Human-Robot Interaction: Characterizing Interface Design in Real-Time Interactive Multimedia Applications.” *ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, 266–273.

Riek, Laurel D., Philip C. Paul, and Peter Robinson. 2010. “When My Robot Smiles at Me: Enabling Human-Robot Rapport via Real-Time Head Gesture Mimicry.” *Journal on Multimodal User Interfaces* 3 (1): 99–108.

Roberts, Jude, and Esther MacCallum-Stewart. 2016. *Gender and Sexuality in Contemporary Popular Fantasy: Beyond Boy Wizards and Kick-Ass Chicks*. Routledge.

Rueben, Matthew, Frank J Bernieri, Cindy M Grimm, and William D Smart. 2017. "Framing Effects on Privacy Concerns about a Home Telepresence Robot." In *Human-Robot Interaction*, 435–444. ACM.

Saakes, Daniel, Vipul Choudhary, Daisuke Sakamoto, Masahiko Inami, and Takeo Igarashi. 2013. "A Teleoperating Interface for Ground Vehicles Using Autonomous Flying Cameras." In *International Conference on Artificial Reality and Telexistence (ICAT)*, 13–19. IEEE.

Sakamoto, Daisuke, Koichiro Honda, Masahiko Inami, and Takeo Igarashi. 2009. "Sketch and Run: A Stroke-Based Interface for Home Robots." *Conference on Human Factors in Computing Systems*, 197–200.

Sakamoto, Daisuke, and Tetsuo Ono. 2006. "Sociality of Robots: Do Robots Construct or Collapse Human Relations?" In *Human-Robot Interaction*, 355. New York, New York, USA: ACM Press.

Sakamoto, Derek (Blizzard). 2015. "Hearthstone: How to Create an Immersive User Interface." In *Game Developers Conference*.

Sanoubari, Elaheh, Stela H. Seo, Diljot Garcha, James E. Young, and Veronica Loureiro-Rodriguez. 2019. "Good Robot Design or Machiavellian? An In-the-Wild Robot Leveraging Minimal Knowledge of Passersby's Culture." In *Human-Robot Interaction*, 382–391. IEEE.

Scarr, Joey, Andy Cockburn, Carl Gutwin, and Andrea Bunt. 2012. "Improving Command Selection with CommandMaps." In *Human Factors in Computing Systems*.

Scarr, Joey, Andy Cockburn, Carl Gutwin, and Philip Quinn. 2011. "Dips and Ceilings." In *Human Factors in Computing Systems*, 2741. New York, New York, USA: ACM Press.

Scheurle, Jennifer. 2017. "Games Aren't Always Fair, the Magic Lies in Making You Think They Are." *Polygon*.

Schoenewolf, Gerald. 1990. "Emotional Contagion: Behavioral Induction in Individuals and Groups." *Modern Psychoanalysis* 15 (1): 49–61.

Schoor, Wolfram. 2006. “VR Based Knowledge Transfer in Medical Technology and Techniques.” In *Computational Science and Its Applications*, 268–277.

Schuemie, Martijn J., Peter van der Straaten, Merel Krijn, and Charles A.P.G. van der Mast. 2001. “Research on Presence in Virtual Reality: A Survey.” *CyberPsychology & Behavior* 4 (2): 183–201.

Seo, Stela H., Denise Geiskkovitch, Masayuki Nakane, Corey King, and James E. Young. 2015a. “Poor Thing! Would You Feel Sorry for a Simulated Robot?” *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI ’15*, 125–132.

Seo, Stela H., Daniel J. Rea, Joel Wiebe, and James E. Young. 2017. “Monocle: Interactive Detail-in-Context Using Two Pan-and-Tilt Cameras to Improve Teleoperation Effectiveness.” In *RO-MAN*. Vol. 2017-Janua.

Seo, Stela H., James E. Young, and Pourang Irani. 2017. “Where Are the Robots? In-Feed Embedded Techniques for Visualizing Robot Team Member Locations.” In *Robot and Human Interactive Communication*, 522–527.

Seo, Stela H, Denise Geiskkovitch, Masayuki Nakane, Corey King, and James E Young. 2015b. “Poor Thing ! Would You Feel Sorry for a Simulated Robot ? A Comparison of Empathy toward a Physical and a Simulated Robot.” In *Human-Robot Interaction*, 125–132.

Sharit, Joseph. 2006. *Handbook of Human Factors and Ergonomics*. Edited by Gavriel Salvendy. Wiley.

Sharma, Megha, Dale Hildebrandt, Gem Newman, James E. Young, and Rasit Eskicioglu. 2013a. “Communicating Affect via Flight Path.” In *Human Robot Interaction 2013*, 293–300.

Sharma, Megha, Dale Hildebrandt, Gem Newman, James E. Young, and Rasit Eskicioglu. 2013b. “Communicating Affect via Flight Path Exploring Use of the Laban Effort System for Designing Affective Locomotion Paths.” *ACM/IEEE International Conference on Human-Robot Interaction*, 293–300.

Sheridan, Thomas B., and R. W. Simpson. 1979. *Toward the Definition and Measurement of the Mental Workload of Transport Pilots*.

Shibata, Takanori, and Kazuyoshi Wada. 2011. "Robot Therapy: A New Approach for Mental Healthcare of the Elderly - A Mini-Review." *Gerontology* 57 (4): 378–386.

Short, E., J. Hart, M. Vu, and B. Scassellati. 2010. "No Fair!! An Interaction with a Cheating Robot." *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 219–226.

Simons, Daniel J., and Christopher F. Chabris. 1999. "Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events." *Perception* 28 (9): 1059–1074.

Singer, Pw. 2009. "Military Robots and the Laws of War." *The New Atlantis* 27: 27–47.

Singh, A, and J E Young. 2012. "Animal-Inspired Human-Robot Interaction: A Robotic Tail for Communicating State." In *Human-Robot Interaction*, 237–238. ACM/IEEE.

Singh, Ashish, Stela H. Seo, Yasmeen Hashish, Masayuki Nakane, James E. Young, and Andrea Bunt. 2013. "An Interface for Remote Robotic Manipulator Control That Reduces Task Load and Fatigue." *RO-MAN*, August. Ieee, 738–743.

Singh, Ashish, and James E. Young. 2013. "A Dog Tail for Utility Robots: Exploring Affective Properties of Tail Movement." *Lecture Notes in Computer Science* 8118 LNCS (PART 2): 403–419.

Skalski, P., R. Tamborini, a. Shelton, M. Buncher, and P. Lindmark. 2011. "Mapping the Road to Fun: Natural Video Game Controllers, Presence, and Game Enjoyment." *New Media & Society* 13 (2): 224–242.

Sloman, Steven A., C. Gordon Hayman, Nobuo Ohta, Janine Law, and et al. 1988. "Forgetting in Primed Fragment Completion." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14 (2): 223–239.

Smitch, Randall C., and Peter Cheeseman. 1986. “Smith, Randall C., and Peter Cheeseman. "On the Representation and Estimation of Spatial Uncertainty.” *Robotics Research* 5 (4): 56–68.

Sollberger, Bernhard, Rolf Rebe, and Doris Eckstein. 2003. “Musical Chords as Affective Priming Context in a Word-Evaluation Task.” *Music Perception* 20 (3): 263–282.

Srinivasan, Raghavan, and Paul P. Jovanis. 1997. “Effect of Selected In-Vehicle Route Guidance Systems on Driver Reaction Times.” *Human Factors: The Journal of the Human Factors and Ergonomics Society* 39 (2): 200–215.

Stainer, Matthew J., Kenneth C. Scott-Brown, and Benjamin W. Tatler. 2013. “Looking for Trouble: A Description of Oculomotor Search Strategies during Live CCTV Operation.” *Frontiers in Human Neuroscience* 7 (September): 615.

Stanton, Neville, Alan Hedge, Karel Brookhuis, Eduardo Salas, and Hal Hendrick. 2006. *Handbook of Human Factors and Ergonomics Methods*. CRC Press.

Steinfeld, A, T Fong, M Field, M Lewis, J Scholtz, and A Schultz. 2006. “Common Metrics for Human-Robot Interaction.” In *Human-Robot Interaction*. ACM/IEEE.

Stenzel, Anna, Eris Chinellato, Maria A. Tirado Bou, Ángel P. del Pobil, Markus Lappe, and Roman Liepelt. 2012. “When Humanoid Robots Become Human-like Interaction Partners: Corepresentation of Robotic Actions.” *Journal of Experimental Psychology: Human Perception and Performance* 38 (5): 1073–1077.

Stoakley, Richard, MJ Conway, and Randy Pausch. 1995. “Virtual Reality on a WIM: Interactive Worlds in Miniature.” *Proceedings of the SIGCHI …*, 265–272.

Stoica, Adrian, Federico Salvioli, and Caitlin Flowers. 2014. “Remote Control of Quadrotor Teams, Using Hand Gestures.” In *ACM/IEEE International Conference on Human-Robot Interaction*, 296–297. New York, New York, USA: ACM Press.

Strauss, Anselm;, and Juliet M. Corbin. 1990. *Basics of Qualitative Research: Grounded Theory Procedures and Techniques.* Sage.

Strauss, Anselm L. 1987. *Qualitative Analysis for Social Scientists*. Cambridge: Cambridge University Press.

Sudnow, David. 1993. *Ways of the Hand: The Organization of Improvised Conduct*. MIT Press.

Sugiura, Yuta, Diasuke Sakamoto, Anusha Withana, Masahiko Inami, and Takeo Igarashi. 2010. "Cooking with Robots." In *Proceedings of the 28th International Conference on Human Factors in Computing Systems - CHI '10*, 2427. New York, New York, USA: ACM Press.

Sung, Ja-Young, Lan Guo, Rebecca E. Grinter, and Henrik I. Christensen. 2007. "'My Roomba Is Rambo': Intimate Home Appliances." *UbiComp 2007: Ubiquitous Computing*, 145–162.

Susilo, Hans. 2015. "Parameter Estimation of Manipulator Robot Dynamics." MIT.

Tanaka, Kazuaki, Naomi Yamashita, Hideyuki Nakanishi, and Hiroshi Ishiguro. 2016. "Teleoperated or Autonomous?: How to Produce a Robot Operator's Pseudo Presence in HRI." In *Human-Robot Interaction*, 133–140. ACM.

Tarasewich, Peter. 2003. "Designing Mobile Commerce Applications." *Communications of the ACM* 46 (12): 57.

Tasse, Dan, Anupriya Ankolekar, and Joshua Hailpern. 2016. "Getting Users' Attention in Web Apps in Likable, Minimally Annoying Ways." *Conference on Human Factors in Computing Systems*, 3324–3334.

Taylor, Publisher, Holly a Yanco, and Jill L Drury. 2009. "Beyond Usability Evaluation : Analysis of Human-Robot Interaction at a Major Robotics Competition." *Robotics*, no. 931414160.

Teng, Wei Chung, Yi Ching Kuo, and Rayi Yanu Tara. 2013. "A Teleoperation System Utilizing Saliency-Based Visual Attention." *IEEE International Conference on Systems, Man, and Cybernetics*, 139–144.

Thellman, Sam, and Tom Ziemke. 2017. "Social Attitudes Toward Robots Are Easily Manipulated." *Proceedings of the Companion of the 2017 ACM/IEEE International*

Conference on Human-Robot Interaction - HRI '17, 299–300.

Thiessen, Raquel, Daniel J. Rea, Diljot S. Garcha, Cheng Cheng, and James E. Young. 2019. "Infrasound for HRI: A Robot Using Low-Frequency Vibrations to Impact How People Perceive Its Actions." In *Human-Robot Interaction*, 11–18. IEEE.

Thornton, Ian M., Heinrich H. Bülthoff, Todd S. Horowitz, Aksel Rynning, and Seong-Whan Lee. 2014. "Interactive Multiple Object Tracking (IMOT)." Edited by Joy J. Geng. *PLoS ONE* 9 (2): e86974.

Torralba, Antonio, Aude Oliva, Monica S Castelhano, and John M Henderson. 2006. "Contextual Guidance of Eye Movements and Attention in Real-World Scenes: The Role of Global Features in Object Search." *Psychological Review* 113 (4): 766–786.

Treue, S., and J. C. Martínez Trujillo. 1999. "Feature-Based Attention Influences Motion Processing Gain in Macaque Visual Cortex." *Nature* 399 (6736): 575–579.

Tsai, Chia-Wen, Pei-Di Shen, and Ya-Ting Fan. 2014. "The Application of Augmented Reality in Online Education:" *International Journal of Information and Communication Technology Education* 10 (2): 75–80.

Tsui, Katherine M, Adam Norton, Daniel J Brooks, Eric McCann, Mikhail S Medvedev, and Holly A Yanco. 2013. "Design and Development of Two Generations of Semi-Autonomous Social Telepresence Robots." In *2013 IEEE Conference on Technologies for Practical Robot Applications (TePRA)*, 1–6. IEEE.

Ullmer, B., and H Ishii. 2000. "Emerging Frameworks for Tangible User Interfaces." *IBM Systems Journal* 39 (3.4): 915–931.

Van Orden, Karl F., Wendy Limbert, Scott Makeig, and Tzyy-Ping Jung. 2006. "Eye Activity Correlates of Workload during a Visuospatial Memory Task." *Human Factors: The Journal of the Human Factors and Ergonomics Society* 43 (1): 111–121.

VanMarle, K., and B. J. Scholl. 2003. "Attentive Tracking of Objects Versus Substances."

Psychological Science 14 (5): 498–504.

Vázquez, Marynel, Aaron Steinfeld, Scott E. Hudson, and Jodi Forlizzi. 2014. "Spatial and Other Social Engagement Cues in a Child-Robot Interaction." In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction - HRI '14*, 391–398.

Veas, Eduardo, Erick Mendez, Steven Feiner, Dieter Schmalstieg, Erick Mendez, Eduardo Veas, Dieter Schmalstieg, and Steve Feiner. 2010. "Directing Attention and Influencing Memory with Visual Saliency Modulation." *ACM Conference on Human Factors in Computing Systems*, 1471–1480.

Walther-franks, Benjamin, Jan Smeddinck, Peter Szmidt, Andrei Haidu, Michael Beetz, and Rainer Malaka. 2015. "Robots , Pancakes , and Computer Games : Designing Serious Games for Robot Imitation Learning." *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, 3623–3632.

Wang, Xu, Diyi Yang, Miaomiao Wen, Kenneth Koedinger, and Carolyn P Rosé. 2015. "Investigating How Student's Cognitive Behavior in MOOC Discussion Forums Affect Learning Gains." In *Educational Data Mining*, 226–233.

Ware, Colin. 2012. *Information Visualization: Perception for Design*. Elsevier.

Warman, Matt. 2011. "Average Video Gamer Is 37." *The Telgraph*, June 8.

Weibel, Roland E, and R John Hansman. 2005. "Safety Considerations for Operation of Unmanned Aerial Vehicles in the National Airspace System." *Transportation* 37 (March): 26–30.

Wentzel, J., D.J. Rea, J.E. Young, and E. Sharlin. 2015. "Shared Presence and Collaboration Using a Co-Located Humanoid Robot." In *HAI 2015 - Proceedings of the 3rd International Conference on Human-Agent Interaction*.

Westlund, Jacqueline M Kory, Marayna Martinez, Maryam Archie, Madhurima Das, and Cynthia Breazeal. 2016. "A Study to Measure the Effect of Framing a Robot as a Social Agent or as a Machine on Children ' s Social Behavior." *Human-Robot Interaction*, 459–460.

Whitworth, Brian. 2005. “Polite Computing.” *Behaviour and Information Technology* 24 (5): 353–363.

Wickens, Christopher D. 2002. “Multiple Resources and Performance Prediction.” *Theoretical Issues in Ergonomics Science* 3 (2): 159–177.

Williams, Kevin W. 2004. “A Summary of Unmanned Aircraft Accident / Incident Data: Human Factors Implications.” *FAA Civil Aerospace Medical Institute*, no. December 2004: 18.

Winkle, Katie. 2019. “Social Influence in HRI with Application to Social Robots for Rehabilitation.” In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 754–756. IEEE.

Yanco, Holly A., and Jill Drury. 2004a. “‘Where Am I?’ Acquiring Situation Awareness Using a Remote Robot Platform.” In *IEEE Systems, Man and Cybernetics*, 3:2835–2840.

Yanco, Holly A., and Jill Drury. 2004b. “Classifying Human-Robot Interaction: An Updated Taxonomy.” *IEEE International Conference on Systems, Man and Cybernetics* 3: 2841–2846.

Yi, Youjae. 1990. “The Effects of Contextual Priming in Print Advertisements.” *Journal of Consumer Research* 17 (2): 215.

Young, J. J., H. Z. Tan, and R. Gray. 2003. “Validity of Haptic Cues and Its Effect on Priming Visual Spatial Attention.” *Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 166–170.

Young, James E. 2010. “Exploring Social Interaction Between Robots and People.” University of Calgary.

Young, James E., Richard Hawkins, Ehud Sharlin, and Takeo Igarashi. 2009. “Toward Acceptable Domestic Robots: Applying Insights from Social Psychology.” *International Journal of Social Robotics* 1 (1): 95–108.

Young, James E., Youichi Kamiyama, Juliane Reichenbach, Takeo Igarashi, and Ehud Sharlin. 2011. "How to Walk a Robot: A Dog-Leash Human-Robot Interface." In *2011 RO-MAN*, 376–382. IEEE.

Young, James E, Ehud Sharlin, and Takeo Igarashi. 2013. "Teaching Robots Style: Designing and Evaluating Style-by-Demonstration for Interactive Robotic Locomotion." *Human–Computer Interaction* 28 (5): 379–416.

Young, James E, Jayoung Sung, Amy Voida, Ehud Sharlin, Takeo Igarashi, Henrik I Christensen, and Rebecca E Grinter. 2010. "Evaluating Human-Robot Interaction." *International Journal of Social Robotics* 3 (1): 53–67.

Zanaty, Peter, Drazen Brscic, and Zsolt Frei. 2008. "3D Visualization for Intelligent Space: Time-Delay Compensation in a Remote Controlled Environment." *Human System Interaction*, 802–807.

APPENDICES

APPENDIX A: STUDY MATERIALS FOR CHAPTER THREE

Below are the questionnaires and recruitment methods used in Chapter Three for investigating how to direct operator attention. The Post-Condition Questionnaire 2 is the NASA Task-Load Index (TLX), from (Hart and Staveland 1988).

Demographics Questionnaire

PARTICIPANT ID: ____

1) What is your age?
   18-20____ 21-25____ 26-30____ 31-35____ 36-40____ 40+____

2) What is your sex?
   Male____ Female____ Intersex____

3) How often do you play 3D videogames such as shooters, racing games…?
   ____Never played videogames
   ____A few times a month or less
   ____Once a week
   ____More than once a week

4) How would you rate your current skill level for this kind of 3D video game?
   Very poor____ Poor____ Fair____ Good____ Very good____

5) How often do you drive a motor vehicle (car, motorcycle, etc.)?
   ____Never driven a vehicle
   ____A few times a month or less
   ____Once a week
   ____More than once a week

6) How would you rate your current vehicle driving skill level?
   Very poor____ Poor____ Fair____ Good____ Very good____

7) How often do you remotely control a vehicle (e.g. car, plane, drone, quadcopter, robot, etc.)?
   ____Never remotely controlled a vehicle
   ____A few times
   ____Every few months
   ____Several times a month

Consent Form

The following poster was printed on The University of Manitoba letterhead

Project Title: Improved interfaces for Robot Tele-operation
Researchers: Dr. James Young, <name removed>, Daniel J. Rea (<emails removed>)

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more details about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

Participation in this study is voluntary, and will take approximately 60 minutes of your time. For this study, you will have the opportunity to interact with a humanoid robot (called NAO) to simulate aspects of search and rescue missions. NAO is a 58 cm tall humanoid (walking) robot. You will use camera feeds to remotely interact with the NAO robot. You will view the video from the robot via a computer. To begin, we will introduce you the NAO robot. We will provide instruction on how we expect you to interact with the robot. You will be given a number of tasks to assess the suitability of the interfaces for the NAO. No expertise or experience is necessary. You will receive $15 for your participation.

All information you provide is considered completely confidential; your name will not be included, or in any other way associated, with the data collected in the study. Data collected during this study will be used for academic research and purpose of publication in an anonymous form. We may use anonymized video or audio data for purposes of public presentation and dissemination only with your express permission (given below). In addition, data will be retained for a maximum of five years in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this study have access. Once published, results of the study will be made available to the public for free at http://home.cs.umanitoba.ca/~young/. Again, no personal information about your involvement will be included. Please note that the University of Manitoba may look at the research records to see that the research is being done in a safe and proper way.

Your signature on this form indicates that you have understood, to your satisfaction, the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from the study at any time, and to refrain from answering any questions asked, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

This research has been approved by the Joint-Faculty Research Ethics Board. If you have any concerns or complaints about this project you may contact Dr. James Young at <removed> or the Human Ethics Secretariat at <removed>. A copy of this consent form has been given to you to keep for your records and reference.

For purposes of research and analysis it is necessary for the experiment to be videotaped.

Do you agree that any video footage taken may also be used for distribution of research, for example, through research videos or images taken from your video? If you say No, your video will be used for internal data analysis purposes only.

No ___ Yes___ but only if you blur my face___ AND/OR if you muffle my voice ___

Participant's Name______________________Signature_________________________ Date________

Researcher's Name______________________Signature_________________________Date________

Post-Condition Questionnaire

PARTICIPANT ID: ____

For the following question, mark **ONE** position along the scale:

Nausea — How queasy, dizzy, or motion sick did you feel?

Very High — Very Low

Enjoyment — How much did you enjoy using the interface?

Not at all — Completely

Awareness — How aware of the robot's surroundings were you during operation?

Not at all Aware — Very Aware

Speed — How quickly were you able to achieve the goals?

Very Slow — Very Quick

Trust — How much did you trust the system?

Did not Trust — Completely Trust

Post-Condition Questionnaire 2

PARTICIPANT ID: ____

For the following question, mark **ONE** position along the scale:

Mental Demand How mentally demanding was the task?

Very Low Very High

Physical Demand How physically demanding was the task?

Very Low Very High

Temporal Demand How hurried or rushed was the pace of the task?

Very Low Very High

Performance How successful were you in accomplishing what you were asked to do?

Perfect Failure

Effort How hard did you have to work to accomplish your level of performance?

Very Low Very High

Frustration How insecure, discouraged, irritated, stressed, and annoyed wereyou?

Very Low Very High

General Experience Questionnaire

PARTICIPANT ID: ____

1) A) Please rank the interfaces from **1** (being the interface that you **MOST** preferred) to 5 (being the interface that you **LEAST** preferred).

(First Design Iteration)

____None ____ Bounce
____Target ____Darken
____Circle

(Second Design Iteration)

____Bounce ____Framed Bounce
____Target ____Tunnel
____Fast Target

The following questions were asked for each interface in the above lists

B) Please describe any pros and cons that you found with each interface.

Bounce

Pros: ____
____

Cons: ____
____

Target

Pros: ____
____

Cons: ____
____

Fast Target

Pros: ____
____

Cons: ____
____

Framed Bounce

Pros: ____
____

Cons: ____
____

Tunnel

Pros: ____

____

Cons: ____

____

2) Did you experience any motion sickness at all? What do you think may have caused or contributed to this?

____

____

3) Do you have any additional positive comments (if any)?

____

____

Do you have any additional negative comments (if any)?

____

____

Do you have any final comments or suggestions (if any)?

____

____

Recruitment Poster

The following poster was printed on The University of Manitoba letterhead

Interact with NAO (a Humanoid Robot) as you complete various tasks in a one-hour human-robot interaction experiment at the University of Manitoba. Note that you must be 18 or over to participate in our experiment.

Please visit:

<Link Removed> or

<Shortlink Removed>

If you have any questions about the study, please contact Daniel J. Rea at <email removed> or Dr. James E. Young at <email and phone number removed>.

This research study was approved by the Joint-Faculty Research Ethics Board, University of Manitoba

APPENDIX B: STUDY MATERIALS FOR EXPLORING PRIMING

Below are the questionnaires and recruitment methods used in Chapter Four for investigating Priming. We first present the materials used for Driving Profile Priming. The materials for Tangible, Descriptive, and No-priming were shared, so we group them together, labelling extra materials for the individual studies when appropriate. The first 6 questions in “Post-condition Questionnaire for Driving Profile Priming” and the “Post-Condition 2 Questionnaire for Tangible, Descriptive, and No Priming” is the NASA Task-Load Index, from (Hart and Staveland 1988).

Post-condition Questionnaire for Driving Profile Priming

1. Mental Demand: How mentally demanding was the task?

Very Low Very High

2. Physical: Demand How physically demanding was the task?

Very Low Very High

3. Temporal Demand: How hurried or rushed was the pace of the task?

Very Low Very High

4. Performance: How successful were you in accomplishing what you were asked to do?

Perfect Failure

5. Effort: How hard did you have to work to accomplish your level of performance?

Very Low Very High

6. Frustration: How insecure, discouraged, irritated, stressed, and annoyed were you?

Very Low Very High

7. Enjoyability: How enjoyable was controlling this robot?

Very enjoyable at all — Not enjoyable

8. Confidency: How confident were you in your ability to avoid hitting obstacles?

Very Unconfident — Very Confident

9. Perception: How heavy did the robot feel?

Very Light — Very Heavy

10. Were there any particular positive points you felt on controlling this robot?

____
____
____
____
____
____

11. Did you have any problems controlling the robot? If yes, please elaborate.

____
____
____
____
____
____

12. Do you have any additional comments about this robot?

____
____
____
____
____
____

Driving Profile Post-Study Questionnaire:

1. Please rank your preference on controlling the robots, with 1 being the robot that you most preferred, and 3 being the robot you least preferred to control.

_____ First robot
_____ Second robot
_____ Third robot

2. Do you have any additional positive comments overall?

____
____
____
____

3. Do you have any additional negative comments overall?

____
____
____
____

4. Do you have any final comments or suggestions?

____
____
____
____

Demographic Questionnaire

1. What is your age?
   Age: ____

2. What is your sex?
   Male____ Female____ Intersex____

How would you rate your current vehicle driving skill level?

Do not drive ____ Very poor____ Poor____ Fair____ Good____ Very good____

Consent Form for Driving Profile Priming

Project Title: Improved interfaces for Robot Teleoperation
Researchers: Dr. James E. Young, <name removed>, Daniel J. Rea, <name removed>, <emails removed>

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more details about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

Participation in this study is voluntary, and will take approximately 60 minutes of your time. For this study, you will have the opportunity to interact with a robot to simulate aspects of search and rescue missions. You will use camera feeds to remotely interact with the robot. You will view the video from the robot via a computer. To begin, we will introduce you to the robot. We will provide instruction on how we expect you to interact with the robot. You will be given a number of tasks to assess the suitability of interfaces for the robot. No expertise or experience is necessary. You will receive up to $15 for your participation.

All information you provide is considered completely confidential; your name will not be included, or in any other way associated, with the data collected in the study. Data collected during this study will be used for academic research and purpose of publication in an anonymous form. We may use anonymized video or audio data for purposes of public presentation and dissemination only with your express permission (given below). In addition, data will be retained for a maximum of five years in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this study have access. Once published, results of the study will be made available to the public for free at http://home.cs.umanitoba.ca/~young/. Again, no personal information about your involvement will be included. Please note that the University of Manitoba may look at the research records to see that the research is being done in a safe and proper way.

Your signature on this form indicates that you have understood, to your satisfaction, the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. **You are free to withdraw from the study at any time**, and to refrain from answering any questions asked, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

This research has been approved by the Joint-Faculty Research Ethics Board. If you have any concerns or complaints about this project you may contact Dr. James Young at < removed> or the Human Ethics Secretariat at <removed>. A copy of this consent form has been given to you to keep for your records and reference.

For purposes of research and analysis it is necessary for the experiment to be videotaped.
Do you agree that any video footage taken may also be used for distribution of research, for example, through research videos or images taken from your video? If you say No, your video will be used for internal data analysis purposes only.

No ___ Yes___ but only if you blur my face___ AND/OR if you muffle my voice ___

Participant's Name____________________ Signature_________________________ Date___________

Researcher's Name_____________________ Signature_________________________ Date__________

Recruitment Poster for Driving Profile Priming

The following poster was printed on The University of Manitoba letterhead

Earn $15 and contribute to science!!

Drive a telepresence robot in a one-hour human-robot interaction experiment at the University of Manitoba. Note that you must be 18 or over to participate in our experiment.

Please visit:

<link removed> or

<shortlink removed>

If you have any questions about the study, please contact Daniel J. Rea at <email removed> or Dr. James E. Young at <email removed> (tel: <removed>).

This research study was approved by the Joint-Faculty Research Ethics Board, University of Manitoba

Study Materials for Tangible, Descriptive, and No Priming

These were filled out on a computer next to the participant

Demographic Questions

* Required

Participant number (do not edit)

Your answer

How old are you? *

Your answer

What gender do you identify as? *

- ◯ Female
- ◯ Male
- ◯ Non-binary
- ◯ Prefer not to say
- ◯ Other: ____

How often do you play 3D video games? *

- ◯ Never
- ◯ A few times a month or less
- ◯ Once a week
- ◯ More than once a week

How often do you drive a motor vehicle? *

- ◯ Never
- ◯ A few times a month or less
- ◯ Once a week
- ◯ More than once a week

How would you rate your current vehicle driving skill level? *

	1	2	3	4	5	
Very Poor	◯	◯	◯	◯	◯	Very Good

NEXT

Post-Condition Questionnaire for Tangible, Descriptive, and No Priming

Telepresence Questions

In general how did you feel the robot performed on the following measures? It may be difficult to answer some questions, but please answer with your best guess about your impressions.

speed *

	1	2	3	4	5	
relatively slow	○	○	○	○	○	relatively fast

weight *

	1	2	3	4	5	
relatively light	○	○	○	○	○	relatively heavy

steering *

	1	2	3	4	5	
relatively uncontrolled	○	○	○	○	○	relatively controlled

durability *

	1	2	3	4	5	
relatively fragile	○	○	○	○	○	relatively durable

power *

	1	2	3	4	5	
relatively weak	○	○	○	○	○	relatively powerful

safety *

	1	2	3	4	5	
relatively dangerous	○	○	○	○	○	relatively safe

responsiveness

	1	2	3	4	5	
relatively unresponsive	○	○	○	○	○	relatively responsive

Post-Condition 2 Questionnaire for Tangible, Descriptive, and No Priming

NASA Task Workload

Please use the scale to choose how you felt about driving the robot

Mental Demand: how mentally demanding was the task? *

	1	2	3	4	5	6	7	8	9	10	
Very Low	○	○	○	○	○	○	○	○	○	○	Very High

Physical Demand: how physically demanding was the task? *

	1	2	3	4	5	6	7	8	9	10	
Very Low	○	○	○	○	○	○	○	○	○	○	Very High

Temporal Demand: how hurried or rushed was the pace of the task? *

	1	2	3	4	5	6	7	8	9	10	
Very Low	○	○	○	○	○	○	○	○	○	○	Very High

Performance: how successful were you in accomplishing the task? *

	1	2	3	4	5	6	7	8	9	10	
Perfect	○	○	○	○	○	○	○	○	○	○	Failure

Effort: how hard did you have to work to accomplish your level of performance?

	1	2	3	4	5	6	7	8	9	10	
Very Low	○	○	○	○	○	○	○	○	○	○	Very High

Frustration: how insecure, discouraged, irritated, stressed, and annoyed were you?

	1	2	3	4	5	6	7	8	9	10	
Very Low	○	○	○	○	○	○	○	○	○	○	Very High

Post-Study Questionnaire for Tangible, Descriptive, and No Priming

General Feedback

Write as much or as little as you want. Any feedback you have that you would like to tell us is useful.

Were there any particular positive points you felt on controlling this robot?

Your answer

Did you have any problems controlling the robot? If yes, please elaborate.

Your answer

Do you have any additional comments about this robot?

Your answer

Consent form for Tangible, Descriptive and No Priming

This consent form was printing on The University of Manitoba letterhead

Project Title: Improved interfaces for Robot Teleoperation
Researchers: Dr. James E. Young, <name removed>, Daniel J. Rea <email removed>
This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more details about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

Participation in this study is voluntary, and will take approximately 60 minutes of your time. For this study, you will have the opportunity to interact with a robot to simulate aspects of search and rescue missions. You will use camera feeds to remotely interact with the robot. You will view the video from the robot via a computer. To begin, we will introduce you to the robot. We will provide instruction on how we expect you to interact with the robot. You will be given a number of tasks to assess the suitability of interfaces for the robot. No expertise or experience is necessary. You will receive up to $15 for your participation.

All information you provide is considered completely confidential; your name will not be included, or in any other way associated, with the data collected in the study. Data collected during this study will be used for academic research and purpose of publication in an anonymous form. We may use anonymized video or audio data for purposes of public presentation and dissemination only with your express permission (given below). In addition, data will be retained for a maximum of five years in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this study have access. Once published, results of the study will be made available to the public for free at http://home.cs.umanitoba.ca/~young/. Again, no personal information about your involvement will be included. Please note that the University of Manitoba may look at the research records to see that the research is being done in a safe and proper way.

Your signature on this form indicates that you have understood, to your satisfaction, the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. **You are free to withdraw from the study at any time**, and to refrain from answering any questions asked, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

This research has been approved by the Joint-Faculty Research Ethics Board. If you have any concerns or complaints about this project you may contact Dr. James Young at <removed> or the Human Ethics Secretariat at <removed>. A copy of this consent form has been given to you to keep for your records and reference.

For purposes of research and analysis it is necessary for the experiment to be videotaped. Do you agree that any video footage taken may also be used for distribution of research, for example, through research videos or images taken from your video? If you say No, your video will be used for internal data analysis purposes only.

No ___ Yes___ but only if you blur my face___ AND/OR if you muffle my voice ___
Participant's Name____________________________ Signature_____________________ Date___________
Researcher's Name____________________________ Signature_____________________ Date___________

Script for Tangible Priming

Thanks for coming today. We are testing 3 different models of telepresence robots. They are all different models of the Double robot, which is an iPad connected to a 2-wheel pendulum robot. They are meant to help people attend conferences, meetings, or family meetings remotely, like Skype on Wheels.

As this type of robot is meant to be used by everyday people, you will be helping us test which is easier to control. Each robot may feel like it drives differently or may interact with our joystick differently, so your feedback will be essential in understanding which robot is better for everyday people to use.

We have the robots in another room. You'll sit here and I'll connect you to the robot from the room, where I'll be able to talk with you. You will run an obstacle course, 3 laps each time, with a different obstacle course for each robot. After each robot trial, I will have you fill out a questionnaire in my web browser.

At any point, feel free to ask questions.

Here is the consent form, it outlines how we will protect your data. Please read it and ask questions, or sign it if you feel you understand and agree with it. <consent form>

Please sign here to show you received your payment. <payment>

Just before we start, please fill out this demographics questionnaire. <demographics>

Now I'll explain how to control the robot. You'll be using this joystick to move the robot, but dont' worry about how complicated it looks, you don't need to hit any buttons. Forward and backwards controls your speed. Side to side will make you turn. You can even turn on the spot. The farther you move the joystick from the centre, the faster it will move in that way.

Does this make sense?

okay, I'll explain the rest of the study in the room. Please put on these headphones.
They will allow you to hear the sound around the robot.

<explain obstacles and laps>
Remember to emphaize:
3 laps,
fast as they can without hitting obstacles.
it's okay to hit things
follow the signs.
They show 90 degree turns.If you see more than 1 arrow, always pick the one that would be in front of you if you turned 90 degrees.

Script Modifications for No Priming

The above script for tangible priming was used as-is for the no priming study with minor modifications. Chiefly, in the second paragraph, instead of describing how the joystick will react with the robots differently, we state "we are simply changing the spring stiffness in the joystick each time. The robot will drive the same way with the same joystick position, just the amount of force you need to move the joystick will change. The robot's abilities never change." Between each condition, we repeat these lines to emphasize the lack of robot change.

Script for Descriptive Priming

Thanks for coming out today.
You will be helping us test 3 different prototypes of a new robot being developed by a robotics company called Double Robotics.
Double Robotics is a leading manufacturer of telepresence robots: robots that someone can control from far away, for example, those that enable everyday people to work from home, to visit international conferences without travelling, to tour areas in other countries, or visit family on the other side of the world. This works by placing a robot in the far away space, putting it on the internet, and letting people connect to it and control it from their own computer. They can see and hear through the robot's camera, and often, the person's face is shown on the robot. See the photo here.

We have been provided with some of their newest prototype robots for testing, and are helping to see how every day, inexperienced people can work with these robots. After all, it is people like you, who are not robotics experts, who will be using this product. So, we are testing various prototypes, and researching which ones are easier or safer to drive. This is why we need you - you are the expert user, who the company hopes will use such a product in the future; while Double Robotics designed the robots in a specific way, they want to know what you think about their potential new products.

We are doing several studies with various robot prototypes. Some prototypes may be very similar and so don't worry if you don't notice any difference. Also, we want your initial reactions, so don't worry too much about being exact and don't over-think your answers.

We will be testing three different robot prototypes today. The three robots are all the same size, and look very similar, but are outfitted with different motors, batteries, wheels, and so on. I'll explain each property (show the three robots now):

Balance: this is a measure of how hard the robot balances as a priority (being safer), versus, moving exactly as you direct it (being more responsive).

Toughness: this is a measure of the material quality used to construct the robot, where tougher robots will not break easily, but may be heavier or less responsive.

Motor power: this is a measure of the power output of the motor, with more powerful motors accelerating and driving faster and more responsively, but poor battery life, and weaker motors being safer and use less energy.

Traction: this is a measure of how well the robot can turn. High traction is safer and easier to use, but costs more to make and uses more battery as it's heavier.

Battery: this is a measure of the battery life of the robot.

The three models we are testing have prototype names (that are kind of cheesy) that encapsulate the design decisions of that robot.

The "Double Turbo" is made to prioritize speed over all else. It is lightweight with a powerful motor, but suffers on toughness. Further, balance, traction, and battery life (smaller battery) are worse to enable this to happen.

The Double Tuff is made to prioritize robustness and product life, and to avoid robot damage and breakage. It has great traction and toughness, but because of the weight is slower (and easier to control), and has less battery life.

The Double Home is a balanced robot with a focus on battery life. It has a great battery, and moderate balance and motor power, so it is not as fast as Double Turbo, but not as robust as Double Tuff.
Each of these ratings is derived from factory specifications. We're interested in how each of these robots

feel to you. To get a sense of driving the robots, you will drive each through an obstacle course. After driving each robot, we will get you to answer questions about your driving experience. In each case, you will get to practice in the obstacle course with the new robot.

explain controls Any questions? *questions*

So first we'll do a test run with an old double model. We will do 3 laps for practice.

Return to room Do you have any questions?

So first you'll be controlling <first condition>. I'll leave the descriptive sheet here while I set up the first course *attach to desk*

go to room, pretend to set up the course for 1 min, connect the robot, explain the obstacle course
Participant completes course.
Return to room, remove robot spec sheet and give questionnaire
Give them the next spec sheet, go to flip the course and "set up next robot"
repeat.

Add final questionnaire:
Give three blank spec sheets (no stars filled) and ask them to fill in the stars according to their experience.
Half stars are allowed.
Underneath each sheet, add a question:
"In what situations do you see this particular robot being useful, and why?"

Bring them to the room, explain deception and why we did it.
Ask for questions
Thank them, it's over.

Post-study Questionnaire Addition for Descriptive Priming

Participants were provided with three of these sheets that were identical except for the robot name. The other two replaced "Double Home" with "Double Turbo" or"Double Tuff"

APPENDIX C: STUDY MATERIALS FOR SOCIAL TELEOPERATION INTERFACES

Below are the questionnaires and recruitment methods used in Chapter Five for investigating how to direct operator attention. The first six questions of "Post-Condition Questionnaire (for Social Teleoperation Interfaces)" is the NASA Task-Load Index (TLX), from (Hart and Staveland 1988). The "Pre-Experiment Questionnaire (for Social Teleoperation Interfaces)" is the Emotional Contagion Questionnaire by (Doherty 1997). The 7-point pictographic scales in the post-condition questionnaire is the Self-Assessment Manikin by (J. Posner et al. 2005).

Demographics Questionnaire

1) What is your age?
   ____

2) What is your sex?
   Male____ Female____ Intersex____

3) How often do you play 3D videogames such as shooters, racing games…?
   ____Never played videogames
   ____A few times a month or less
   ____Once a week
   ____More than once a week

4) How would you rate your current skill level for this kind of 3D video game?
   Very poor____ Poor____ Fair____ Good____ Very good____

5) How would you rate your current vehicle driving skill level?
   Don't drive____ Very poor____ Poor____ Fair____ Good____ Very good____

Have you participated in our studies before?

____

Consent Form for Social Teleoperation Interfaces

Project Title: Improved interfaces for Robot Teleoperation
Researchers: Dr. James E. Young, <name removed>, Daniel J. Rea <emails removed>

This consent form, a copy of which will be left with you for your records and reference, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more details about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

Participation in this study is voluntary, and will take approximately 60 minutes of your time. For this study, you will have the opportunity to interact with a robot to simulate aspects of search and rescue missions. You will use camera feeds to remotely interact with the robot. You will view the video from the robot via a computer. To begin, we will introduce you to the robot. We will provide instruction on how we expect you to interact with the robot. You will be given a number of tasks to assess the suitability of interfaces for the robot. No expertise or experience is necessary. You will receive up to $15 for your participation.

All information you provide is considered completely confidential; your name will not be included, or in any other way associated, with the data collected in the study. Data collected during this study will be used for academic research and purpose of publication in an anonymous form. We may use anonymized video or audio data for purposes of public presentation and dissemination only with your express permission (given below). In addition, data will be retained for a maximum of five years in a locked office in the EITC building, University of Manitoba, to which only researchers associated with this study have access. Once published, results of the study will be made available to the public for free at http://home.cs.umanitoba.ca/~young/. Again, no personal information about your involvement will be included. Please note that the University of Manitoba may look at the research records to see that the research is being done in a safe and proper way.

Your signature on this form indicates that you have understood, to your satisfaction, the information regarding participation in the research project and agree to participate as a subject. By doing this you also confirm that you are of the age of majority in Canada (18 years or more). In no way does this waive your legal rights nor release the researchers, sponsors, or involved institutions from their legal and professional responsibilities. **You are free to withdraw from the study at any time**, and to refrain from answering any questions asked, without prejudice or consequence. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

This research has been approved by the Joint-Faculty Research Ethics Board. If you have any concerns or complaints about this project you may contact Dr. James Young at <removed> or the Human Ethics Secretariat at <removed>. A copy of this consent form has been given to you to keep for your records and reference.

For purposes of research and analysis it is necessary for the experiment to be videotaped.

Do you agree that any video footage taken may also be used for distribution of research, for example, through research videos or images taken from your video? If you say No, your video will be used for internal data analysis purposes only.

No ___ Yes___ but only if you blur my face___ AND/OR if you muffle my voice ___

Participant's Name__________________________ Signature_____________________ Date_________________
Researcher's Name_________________________ Signature_____________________ Date_________________

Pre-Experiment Questionnaire (for Social Teleoperation Interfaces)

For the following question, mark **ONE** position along the scale:

1) If someone I'm talking with begins to cry, I get teary-eyed

Never	Rarely	Usually	Often	Always
○	○	○	○	○

2) Being with a happy person picks me up when I'm feeling down.

Never	Rarely	Usually	Often	Always
○	○	○	○	○

3) When someone smiles warmly at me, 1 smile back and feel warm inside.

Never	Rarely	Usually	Often	Always
○	○	○	○	○

4) 1 get filled with sorrow when people talk about the death of their loved ones.

Never	Rarely	Usually	Often	Always
○	○	○	○	○

5) 1 clench my jaws and my shoulders get tight when 1 see the angry faces on the news.

Never	Rarely	Usually	Often	Always
○	○	○	○	○

6) When 1 look into the eyes of the one I love, my mind is filled with thoughts of romance.

Never	Rarely	Usually	Often	Always
○	○	○	○	○

7) It irritates me to be around angry people

Never	Rarely	Usually	Often	Always
○	○	○	○	○

8) Watching the fearful faces of victims on the news makes me try to imagine how they might be feeling.

Never Rarely Usually Often Always

9) I melt when the one I love holds me close.

Never Rarely Usually Often Always

10) I tense when overhearing an angry quarrel.

Never Rarely Usually Often Always

11) Being around happy people fills my mind with happy thoughts

Never Rarely Usually Often Always

12) I sense my body responding when the one I love touches me

Never Rarely Usually Often Always

13) I notice myself getting tense when I'm around people who are stressed out.

Never Rarely Usually Often Always

14) I cry at sad movies.

Never Rarely Usually Often Always

15) Listening to the shrill screams of a terrified child in a dentist's waiting room makes me feel nervous.

Never Rarely Usually Often Always

Post-Condition Questionnaire (for Social Teleoperation Interfaces)

For the following question, mark **ONE** position along the scale:

Mental Demand How mentally demanding was the task?

Very Low — Very High

Physical Demand How physically demanding was the task?

Very Low — Very High

Temporal Demand How hurried or rushed was the pace of the task?

Very Low — Very High

Performance How successful were you in accomplishing what you were asked to do?

Perfect — Failure

Effort How hard did you have to work to accomplish your level of performance?

Very Low — Very High

Frustration How insecure, discouraged, irritated, stressed, and annoyed wereyou?

Very Low — Very High

How safe is this robot for driving?

Not safe at all Very safe

How informative was the interface?

Not informative Very Informative

Please circle the picture that best represents how you feel right now. Select one item from each set

Negative Positive

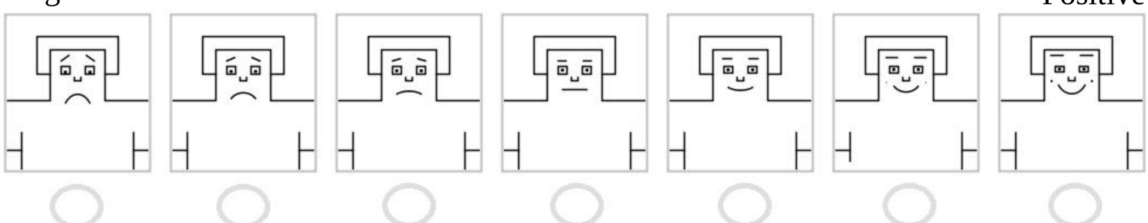

Low High

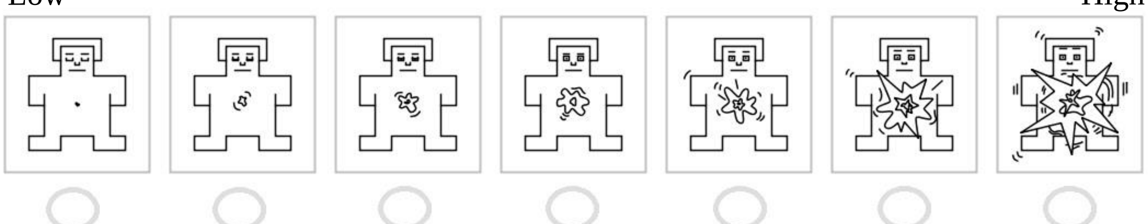

1) Do you have any additional positive comments (if any)?

____

____

____

2) Do you have any additional negative comments (if any)?

____

____

____

3) Do you have any final comments or suggestions (if any)?

____

____

____

Post-Condition Questionnaire 2 (for Social Teleoperation Interfaces)

Which expression was the interface closest to *most often*?

    

How often did you have a collision?

Never | Occasionally | Sometimes | Frequently | Constantly

    

How often did you have a collision?

Never | Occasionally | Sometimes | Frequently | Constantly

What was your *average velocity* closest to?

20% | 40% | 60% | 80% | 100%

How often did you have a collision?

Never | Occasionally | Sometimes | Frequently | Constantly

General Experience Questionnaire (for Social Teleoperation Interfaces)

Please rank each interface in order of preference, 1 being your favourite

___ First face interface
___ Second face interface
___ Numeric interface

1) Do you have any additional positive comments (if any)?
____
____
____

2) Do you have any additional negative comments (if any)?
____
____
____

3) Do you have any final comments or suggestions (if any)?
____
____
____

We are interested in repeating this study in the future to investigate long-term effects. These follow-up studies will be similar to this one, and you will receive 15$ compensation again if you participate. If you would like, please provide an email to contact you for this opportunity:

(optional) Contact Email: ______________________________
Your email will not be shared with any other party.

APPENDIX D: LIST OF VIDEO GAME REFERENCES

Games in italics were NOT used in the video game survey (Chapter Six). Games with an * in the survey were played, not watched.

Tetris, Pajitnov, A.L., 1984.

The Legend of Zelda Series, Nintendo, 1986-2018.

*Final Fantasy Series, Square Enix, 1987-2018.

Doom, id Software, 1993.

Hi-Octane, Bullfrog Productions, 1995.

*Pokemon (Series), Game Freak, 1996-2019

*Goldeneye 007, Rare, 1997.

Gran Turismo, Sony, 1997.

Tekken 3, Namco, 1998.

*The Legend of Zelda: Ocarina of Time, Nintendo, 1998.

NFL2K, Sega, 1999.

*Baldur's Gate II, BioWare, 2000.

*Tony Hawk's Pro Skater 2, Neversoft, 2000.

Hey You, Pikachu!, Nintendo, 2000.

*Perfect Dark, Rare, 2000.

Halo: Combat Evolved, Bungie, 2001.

Metal Gear Solid 2: Sons of Liberty, Konami, 2001.

*The Legend of Zelda: Wind Waker, Nintendo, 2002.

MechWarrior 4: Mercenaries, Microsoft, 2002

*Kingdom Hearts Series, Square Enix, 2002-2019.

*Metroid Prime, Retro Studios, 2002.

*Soul Calibur 2, Namco Bandai, 2002.

*World of Warcraft, Blizzard, 2004.

*Half-life 2, Valve, 2004.

Resident Evil 4, Capcom, 2005.

*The Legend of Zelda: Twilight Princess, Nintendo, 2006

*Super Mario Galaxy, Nintendo, 2007.

*Bioshock, 2K Games, 2007.

*Assassin's Creed, Ubisoft, 2007.

Team Fortress 2, Valve, 2007

GTA IV, Rockstar Games, 2008.

LittleBigPlanet, Sony, 2008.

Uncharted 2, Naughty Dog, 2009.

Left 4 Dead 2, Valve, 2009.

*Dark Souls Series, From Software, 2009-2016

*Mass Effect 2, BioWare, 2010.

Red Dead Redemption, Rockstar Games, 2010.

*Starcraft 2, Blizzard, 2010.

Mario Kart 7, Nintendo, 2011.

*Batman: Arkham City, Rocksteady Studios, 2011.

*Skyrim, Bethesda, 2011.

Mass Effect 3, BioWare, 2012.

Spec-Ops: The Line, Yager Development, 2012.

Kinect Star Wars, LucasArts, 2012.

The Last of Us, Naughty Dog, 2013.

DmC: Devil May Cry, Ninja Theory, 2013.

*Final Fantasy XIV, Square-Enix, 2013.

DragonAge: Inquisition, BioWare, 2014.

*Overwatch, Blizzard, 2016.

Rocket League, Psyonix, 2016.

*Destiny 2, Bungie, 2017.

*The Legend of Zelda: Breath of the Wild, Nintendo, 2017.

APPENDIX E: ETHICS APPROVAL FORMS

All of the research in this thesis was approved by the University of Manitoba Joint-Faculty Research Ethics Board. We include the certificates of ethics approval for our experiments, as well as proof of the required ethics training below.

TCPS 2: CORE Certificate

The University of Manitoba requires all researchers working with people to complete an ethics training course. Below is proof that I have completed this requirement.

PANEL ON RESEARCH ETHICS
Navigating the ethics of human research

TCPS 2: CORE

Certificate of Completion

This document certifies that

Daniel Rea

has completed the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans Course on Research Ethics (TCPS 2: CORE)

Date of Issue: **4 October, 2011**

Ethics Approval for Experiments

We received ethics approval for each study in this thesis. Due to the similarities of the experiments, we submitted amendments to our original ethics form instead of a brand new study procedure. Thus, the first copy is of our original ethics approval, and all following documents are the approval to update our ethics with our new experiment protocol.

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics
[redacted]

APPROVAL CERTIFICATE

April 1, 2015

NSERC

TO: **James E. Young**
Principal Investigator [redacted]

FROM: **Susan Frohlick, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:020**
"Exploring alternate camera systems for remotely controlling robots"

Please be advised that your above-referenced protocol has received human ethics approval by the **Joint-Faculty Research Ethics Board**, which is organized and operates according to the Tri-Council Policy Statement (2). **This approval is valid for one year only**.

Any significant changes of the protocol and/or informed consent form should be reported to the Human Ethics Secretariat in advance of implementation of such changes.

Please note:

- **If you have funds pending human ethics approval, please mail/e-mail/fax** [redacted] **a copy of this Approval (identifying the related UM Project Number) to the Research Grants Officer in ORS in order to initiate fund setup. (How to find your UM Project Number: http://umanitoba.ca/research/ors/mrt-faq.html#pr0)**

- **if you have received multi-year funding for this research, responsibility lies with you to apply for and obtain Renewal Approval at the expiry of the initial one-year approval; otherwise the account will be locked.**

The Research Quality Management Office may request to review research documentation from this project to demonstrate compliance with this approved protocol and the University of Manitoba *Ethics of Research Involving Humans*.

The Research Ethics Board requests a final report for your study (available at: http://umanitoba.ca/research/orec/ethics/human_ethics_REB_forms_guidelines.html**) in order to be in compliance with Tri-Council Guidelines.**

umanitoba.ca/research

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics
[redacted]

AMENDMENT APPROVAL

February 29, 2016

TO: **James E. Young** [redacted]
Principal Investigator

FROM: **Lorna Guse, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:020**
"Exploring Alternate Camera Systems for Remotely Controlled Robots"

This will acknowledge your Amendment Request dated February 29, 2016 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Coordinator in advance of implementation.

umanitoba.ca/research

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics
[redacted]

AMENDMENT APPROVAL

June 9, 2016

TO: **James E. Young** [redacted]
Principal Investigator

FROM: **Lorna Guse, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:020 (HS17568)**
"Exploring Alternate Camera Systems for Remotely Controlled Robots"

This will acknowledge your Amendment Request dated June 7, 2016 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Coordinator in advance of implementation.

umanitoba.ca/research

UNIVERSITY OF MANITOBA | Research Ethics and Compliance
Office of the Vice-President (Research and International)

Human Ethics
[redacted]

AMENDMENT APPROVAL

August 26, 2016

TO: **James E. Young** [redacted]
Principal Investigator

FROM: **Lorna Guse, Chair**
Joint-Faculty Research Ethics Board (JFREB)

Re: **Protocol #J2015:020 (HS17568)**
"Exploring Alternate Camera Systems for Remotely Controlled Robots"

This will acknowledge your Amendment Request dated August 26, 2016 requesting amendment to your above-noted protocol.

Approval is given for this amendment. Any further changes to the protocol must be reported to the Human Ethics Coordinator in advance of implementation.

umanitoba.ca/research

UNIVERSITY OF MANITOBA | Research Ethics and Compliance

Human Ethics
[redacted]

RENEWAL APPROVAL

Date: March 26, 2018 **New Expiry:** March 31, 2019

TO: **James E. Young**
Principal Investigator

FROM: **Kevin Russell, Chair**
Joint-Faculty Research Ethics Board (JFREB)

[redacted]

Re: **Protocol #J2015:020 (HS17568)**
"Exploring Alternate Camera Systems for Remotely Controlled Robots"

Joint-Faculty Research Ethics Board (JFREB) has reviewed and renewed the above research. JFREB is constituted and operates in accordance with the current *Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans.*

This approval is subject to the following conditions:

1. Any modification to the research must be submitted to JFREB for approval before implementation.
2. Any deviations to the research or adverse events must be submitted to JFREB as soon as possible.
3. This renewal is valid for one year only and a Renewal Request must be submitted and approved by the above expiry date.
4. A Study Closure form must be submitted to JFREB when the research is complete or terminated.

Funded Protocols:
- **Please mail/e-mail a copy of this Renewal Approval, identifying the related UM Project Number, to the Research Grants Officer in ORS.**

Research Ethics and Compliance is a part of the Office of the Vice-President (Research and International)
umanitoba.ca/research