THESIS

THE IMPACT OF MANIPULATIVE CONTENT ON HUMAN PERFORMANCE IN AUGMENTED REALITY

Submitted by

Evan D. Anspach

Department of Computer Science

In partial fulfillment of the requirements

For the Degree of Master of Science

Colorado State University

Fort Collins, Colorado

Summer 2025

Committee:

Advisor: Indrakshi Ray

Mohammed Sayafet Arefin
Rosa Martey

ABSTRACT

THE IMPACT OF MANIPULATIVE CONTENT ON HUMAN PERFORMANCE IN AUGMENTED REALITY

Extended Reality is the spectrum of spaces and experiences, both virtual and augmented which include both Augmented Reality (AR) and Virtual Reality (VR). Of the two categories, Optical See-Through (OST) Augmented Reality is beginning to be used more widely in the public domain. However, addressing manipulative content is necessary for the widespread adoption of OST AR technology. Extended Reality (XR) Devices have had many vulnerabilities identified in previous works that may make them susceptible to the introduction of manipulative content, which an attacker may be able to use for a variety of purposes. For instance, in a cybersecurity context, attackers might try to influence and reduce user performance by changing the quality of AR information, introducing misleading content, irrelevant data, and other adverse factors. This may allow the attackers to control user behavior, slow down or stop important tasks performed in XR or to annoy or otherwise adversely affect the mental state of the XR user. This research investigates how helpful, misleading, and irrelevant information in OST AR affects human performance. The study used a memory task and employed a repeated measures design involving 19 participants. The findings revealed that the participants needed more time to complete the task when presented with irrelevant information compared to when they had access to useful AR information or when AR content was not presented. In addition, helpful AR information allowed users to complete the task more effectively with fewer errors than irrelevant and misleading AR information. The results suggest that AR enhances user memory, enabling them to perform tasks more efficiently. Moreover, when malicious information is introduced, manipulative content can effectively increase the decision-making time of their targets by disrupting memory-based judgments.

ACKNOWLEDGEMENTS

I first acknowledge and thank my coauthors on the published work "The Impact of Relevant Augmented Reality Information on Human Performance" [20] which makes up the majority of this thesis, Matthew Sturgeon, Dr. Mohammed Sayafet Arefin, Dr. Indrakshi Ray and Dr. Francisco Ortega. Without their hard work and dedication this work would not have been possible.

I express my greatest gratitude to my advisor Professor Indrakshi Ray for her advice, support, encouragement, and guidance throughout my Master's program. Professor Ray's insights significantly impacted and contributed to the success of my work and overall academic career. Likewise I express my gratitude to Professor Mohammed Sayafet Arefin for his mentorship and guidance throughout my work with his lab.

I thank Professor Rosa Mikeal Martey for their time in serving in my Master's advisory committee. I thank all of the professors and instructors who have inspired me though their teaching and conversation throughout my academic career.

I thank my friends and colleagues both in the DBSec group as well as in the ARefin Lab for their unending support and encouragement which pushed me forward to achieve this goal.

Finally, I thank the funding agencies for supporting this work: This work was partially supported by the U.S. National Science Foundation under Grant No. DMS 2123761, CNS 1822118, AMI, NewPush, Cyber Risk Research, ARL, NIST under Award No. 60NANB23D152, the State of Colorado (grant #SB 18-086), and Colorado State University internal funds.

DEDICATION

I dedicate this thesis to my partner, Ari - who's patience and understanding has been pivotal to my work. You are my solid rock to stand on when the seas rise.

TABLE OF CONTENTS

Chapter 1

Introduction

1.1 Motivation

Optical See-Through (OST) Augmented Reality (AR) technology overlays virtual information onto the real-world environment, thereby improving the comprehension of the surroundings. Owing to the accessibility of commercially available OST AR devices (such as Microsoft HoloLens 2, Magic Leap 2, and others), both investors and commercial enterprises are showing a growing interest in offering AR technologies for corporate use as well as for general public adoption. However, this increased interest raises concerns about how the quality of AR information presented effects the user's performance in the headset. Malicious actors who have begun to view these devices not as the future of technology-enhanced realities but rather as a new target for exploitation to achieve nefarious ends may want to take advantage of the impact that information quality has on a user to manipulate them. Indeed, vulnerabilities have already been discovered in almost all parts of these devices' systems, from the hardware to the applications running on top of them. Manipulating the information presented on these objects may impact the user's ability to utilize the AR information or to interact correctly with the real objects, leading users to make incorrect decisions or impairing user task performance when the AR is present. This offsets the intended use of these devices; instead of providing useful information and enhancing user interaction, the virtual becomes a hindrance instead. An example of this could be a cyber-attacker feeding incorrect details into a navigation display, causing a user to take a wrong turn and fail to reach their intended destination promptly, or guiding them to a particular place.

1.2 Description

In this thesis, we focus on two categories of *manipulative content*, misleading content and irrelevant content. *Misleading content* refers to false information and aims to persuade a user to make

an erroneous decision based on it. *Irrelevant content* is content presented in the object that is not relevant to the task the user is trying to achieve. The primary objective of our research is to explore how manipulated virtual content affects human performance when using the OST AR system. To accomplish this, we performed a memory-based task using the card game "Concentration" with an assistive OST AR display. The OST AR display's content was either manipulative or helpful. AR *helpful content* provides accurate information to the user about the task. To the best of our knowledge, this research is the first effort to investigate the impact of manipulative content on human memory knowledge within an AR system.

1.3 Research Questions

This work was motivated by the three main research questions below.

RQ1: What is the effect of *helpful AR* information on human performance?

RQ2: What is the effect of manipulative (*misleading or irrelevant AR*) information on human performance?

RQ3: What is the effect of *helpful, misleading and irrelevant AR* when presented in a combined form on human performance?

Based on the above research questions we developed the following four hypotheses for our research.

1.4 Hypotheses

H1: Helpful AR information would increase human performance.

H2: Misleading AR information would decrease human performance when presented during a memory task.

H3: Irrelevant AR information would decrease human performance when presented during a memory task.

H4: A combination of misleading, irrelevant, and helpful AR information would yield human performance similar to that observed when only either misleading or irrelevant information is provided.

Investigating these hypotheses enables us to make some inferences about how users are affected by manipulative content, their reactions to it, and what strategies they employ to mitigate its effects. This provides a better understanding of the effectiveness of this form of manipulation on users and highlights some related issues in this area that need further investigation.

1.5 Contributions

My contributions to this work included design of the experiment, co-analysis of participant data, lead in drafting, providing direction and feedback on application design, writing and editing the manuscript, and responding to and implementing reviewer feedback.

1.6 Outline

The following thesis is laid out as follows. First, we introduce a summary of other contributions within the field of AR cybersecurity, then we discuss the design and methods used in our experiment, we then analyze the quantitative and qualitative data of the participants and discuss the results. We conclude with an overview of the work and future research directions that can be taken in this space.

Chapter 2

Literature Review [1]

The ultimate goal of the OST AR system is to provide accurate and helpful information, allowing users to perform their tasks successfully. Due to this ability, AR information has been extensively utilized in various fields, including AR-based medical surgery [2], AR for indoor and outdoor navigation [14], AR manufacturing guidance [6], AR in education [3], among many others. Although AR has been suggested as a beneficial technique or resolution in several fields, some of these fields might experience fewer disastrous consequences if AR information were manipulated than others. For example, one could presume that the disastrous impact on AR-enhanced museums [7] would be less severe than in AR-based military [13], air traffic control [9] or surgery [18].

2.1 Cybersecurity Threats to perception in Extended Reality

XR devices, due to both being computing devices and having virtual content introduced in an immersive nature are vulnerable to both traditional computing attacks as well as more novel Perceptual attacks. Traditional atatcks such as network attacks, attacks on the device's resources such as denial of service or framerate overload attacks, or data leakage attacks that cause privacy concerns often mirror their traditional computing counterparts. Perceptual Attacks or Perception Manipulation Attacks (PMAs) seek to leverage the immersive nature of the device to impact some part of the human perceptual process through some access to the device.

2.1.1 Side Channel Attacks and Privacy

XR devices have significant concerns related to privacy and disclosure of information, particularly due to side channel attacks allowing attackers to infer information about what activities an

[1] Sturgeon, M., Anspach, E., Ortega, F., Ray, I., Safayet Arefin, M.: Impact of Relevant Augmented Reality Information on Human Performance. In: Bebis, G., Patel, V., Gu, J., Panetta, J., Gingold, Y., Johnsen, K., Arefin, M.S., Dutta, S., Biswas, A. (eds.) Ad- vances in Visual Computing. pp. 185–198. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-77389-1_15

user may be engaged with in the virtual space. Ling et al. [12] investigated how user actions in a VR environment can be observed by an outside observer using a camera recording system or through compromise of the device's sensor systems in order to determine patters on user keystrokes. This work found that with reasonable certainty an attacker can determine which keystrokes a user has entered through observation of the sensors or the subject in order to learn confidential information such as passwords, financial information or personal information that is entered through a VR keyboard. Arafat et al. [1] investigated how the channel state information of Wifi signals can be intercepted and interpreted to recognize virtual keystrokes in VR headsets. They found a 65.75% accuracy in recognizing which keystrokes were inputted. The most notable part of their work is the ease of setup of their approach as it does not require advanced models with access to the device or special sensors such as camera's near the victim, this makes the attack relatively easy to deploy despite its lower accuracy than some other approaches. Shi et al. [17] investigated the privacy and data concerns related to how the AR/VR headsets are closely mounted on the face, thus encouraging privacy leakage through the voice interfaces and motion sensors of the device. They designed an eavesdropping attack which uses signal source separation to determine biometric information about the user such as gender, the speakers identity, or the content of the user's speech.

2.1.2 Network Attacks

Other works have identified network threats to XR devices that may impact user perception and performance inside the system. Gulhane et al. [8] investigated how network attacks such as packet loss, and packet sniffing impact Virtual Reality Learning Environments (VRLEs). They found significant privacy concerns due to packet sniffing in VR applications and demonstrated that confidential information about the VR user and their application can be captured and deciphered in a malicious environment. They also found that session failures and network discrepancies caused through packet loss which could be caused by attacks such as denial of service and other network flooding attacks severely impact the user's ability to access content in the learning environment

causing confusion or frustration. Impacts to latency in network based applications can also cause motion sickness and other negative physiological responses to users.

2.1.3 Perception Manipulation

Several previous works have studied many different techniques to achieve manipulations of different parts of human perception, many of which may be relevant to information quality in an OST AR scenario. Casey et al. [4] described a system to induce physical movement of a user, cause physical collisions to a user, disorient and cause motion sickness to the user, and create overlays to distract or cause cognitive distress to the user in a virtual reality (VR) system. Their mentioned "Overlay Attack" allowed an attacker to create custom virtual overlays which are rendered in the device view-space and cannot be removed easily by the user. These attacks exploit vulnerabilities in the inter-process communication binder kernel of the underlying Android operating system of the device, which allows the introduction of arbitrary information into the virtual environment. Of particular note is their method of remotely disabling safeguards in the system such as the VR boundary that outlines the allowed movement spaces for the VR session, using techniques like theirs safeguards implements in apps on the headset could be bypassed by attackers to achieve their goals.

Casey's work in the human joystick attack built off of the work that was done earlier by Sun et al. In their study of redirected walking and infinite walking [21]. Although this mechanism is generally benign in nature and intended to be a positive enhancement to immersive experiences, it can be used maliciously to disconnect a user's position in virtual space from their position in physical space; causing users to be unaware of where they actually are in physicality. This can be used by an attacker to direct an immersed user to a specific location that will cause distress or physical harm to the user or to steer the user to a location that gives the attacker some access to the user physical or through the network.

Tseng et al. [22] described various harmful situations that could arise from perceptual manipulations. They identified Perception Manipulation Attacks (PMAs) as a significant threat to AR

interactions. They describe a variety of attack classifications, but of particular importance to our work is the mismatching attack, which involves a virtual object being incorrectly aligned with its counterpart, leading to user misinterpretation. In this attack the virtual object's position is shifted to where it no longer reflects the accurate position of a matching physical object. This mismatch will cause user interactions that are meant for the physical object to fail, such as sitting on a chair that is present in both a virtual and a physical environment. In this case the user may fall and suffer distress or injury. A possible concern for this kind of attack is in AR plumbing, medical or military uses, in these technologies misalignment of virtual and real objects could cause significant damages or deaths.

Earlier works have also worked with forms of perception manipulation, Punponsanon et al in two different works studied how XR output can affect human perception of haptic softness and bending stiffness [15, 16]. These works indicated that XR visual stimuli can influence haptic perception of objects. This suggests that an attacker could manipulate the visual image in the XR environment to influence a user to make incorrect decisions about an object due to conflicts between the visual stimuli and the haptic response.

Cheng et al. [5] introduced their own attack methods that interacted with visual, auditory, and situational perceptions. Their work designed three studies, each directed a different PMA at the user's perception of the task. They found that introduction of the PMA's decreased situational awareness, slowed reaction times, and induced incorrect decisions. They found that participants were impacted by the malicious content and that it often resulted in changes to the mental models of the users around the interactions with the XR device and environment. Previous studies have demonstrated that altered AR overlay content can affect users' perception of elements. However, no previous research has investigated how individuals manage task performance considering their memory knowledge when confronted with manipulative content.

Chapter 3

Study Methodology [2]

In this chapter we lay out the design and implementation of the study performed including descriptions of the Apparatus on which the experiment was performed, the task that was performed for the experiment, the procedure of the study, and information regarding the participants and their demographics.

3.1 Apparatus and Setup

The experiment used the Microsoft HoloLens 2, a binocular see-through AR display. A standalone Unity application (version 2022.3.23f1) managed the entire experiment. This application utilized the Microsoft Mixed Reality Toolkit (MRTK) and was executed on an MSI laptop running Windows 10. A deck of bicycle playing cards was employed for the physical cards in the study. The cards were laid out on a table draped with a black cloth to reduce reflections and enhance contrast and legibility. The cards had four possible suits: Diamonds, Hearts, Clubs, and Spades. We considered a pseudo-random grouping of card values with the red 5s, 8s, 9s, and Queens and the black 2s, 6s, 7s, and Jacks. Within those suits there were two colors, Red (Hearts, Diamonds) and Black (Spades, Clubs). Thus, 16 of the 52 cards were used in a 4x4 grid arrangement. An individual QR code was affixed to the back of each card with tape. QR codes were used to produce AR information related to the suit and value of the card. AR content appeared on the right side of the card as an overlay to ensure that it did not interfere with the physical card's content. An instruction sheet for the training rounds was placed adjacent to the card grid.

[2] Sturgeon, M., Anspach, E., Ortega, F., Ray, I., Safayet Arefin, M.: Impact of Relevant Augmented Reality Information on Human Performance. In: Bebis, G., Patel, V., Gu, J., Panetta, J., Gingold, Y., Johnsen, K., Arefin, M.S., Dutta, S., Biswas, A. (eds.) Ad- vances in Visual Computing. pp. 185–198. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-77389-1_15

3.2 Experimental Task

We considered *Card Game Concentration* as our experimental task. This memory game requires the player to recall specific information. The flow of a round of the experimental task was as follows.

Step 1: The participant observed that there were 16 cards placed face down.

Step 2: The participant turned two cards face up during each turn. Each turn, the participant tried to find pairs of cards with matching values and colors based on their suits. For instance, if the ace of hearts and the ace of spades were revealed, it would result in a red ace and a black ace, making it an invalid match. Conversely, if the ace of clubs and the ace of spades were turned over, both being black aces would result in a valid match.

Step 3: If the participant determined that the colors and values of both cards were identical, they removed the cards from the game space. Otherwise, the cards were turned back over and remained in the game space.

Step 4: Continue executing Steps 2 and 3 until no cards remain in the game space.

As participants attempted to finish the task, AR content was displayed via the Hololens 2 OST AR display. The overlay data became visible only after a card was revealed by the participant for the first time, and remained revealed for the remainder of the trial. This was managed by an experimenter using a bluetooth keyboard.

3.3 Experimental Variables and Design

We employed one independent variable referred to as *Types of AR content*. This variable consisted of five distinct levels:

No AR: In this level, participants interacted with the task while wearing the Hololens 2, but no virtual content was displayed (see Fig. 3.1b).

AR Helpful: When the AR content was shown, it accurately conveyed the true information of the physical card (see Fig. 3.1c).

AR Irrelevant: When AR content was displayed, it provided unrelated information of the physical card (see Fig. 3.1d).

AR Misleading: When AR content was displayed, it provided incorrect information of the physical card (see Fig. 3.1e).

AR Combined: When AR content was shown, it offered information about the physical card that could be helpful, misleading, or irrelevant (refer to Fig. 3.1f).

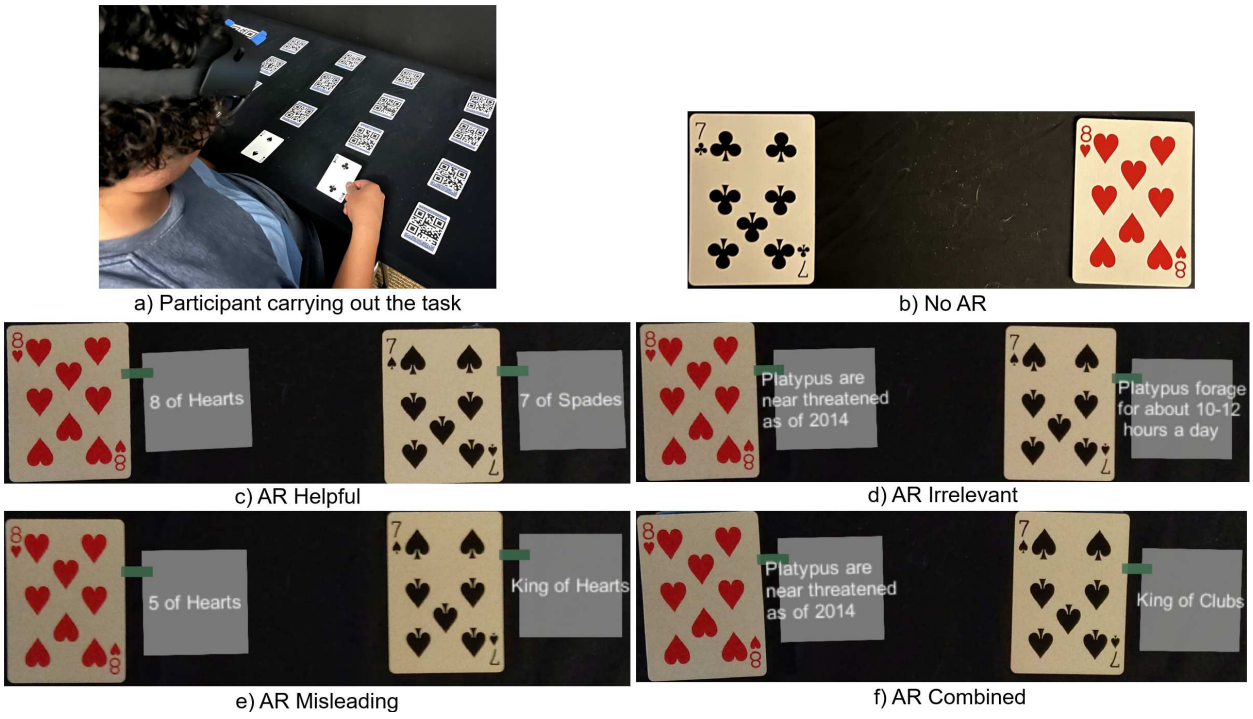

Figure 3.1: (a) Participant performed a card game concentration as experimental task. (b)-(f) Participant's view of different AR conditions as seen through the Microsoft HoloLens 2. (b) In the No AR condition, there was no AR overlay shown. (c) AR Helpful: AR display provided accurate descriptions of the cards to the participant. (d) AR Irrelevant: AR display presents details about platypuses, which are unrelated and offer no useful information for the task. (e) AR Misleading: AR display information inaccurately represents the cards; it indicates that the eight of hearts is the five of hearts and the seven of spades is the king of hearts. (f) AR Combined: AR display provides information that is irrelevant for one card in the form of a platypus fact, and is misleading for the other card in that it incorrectly states that the seven of spades is the king of clubs. The combined condition also included the helpful condition during trials alongside the irrelevant and misleading information.

We measured three dependent variables in the experiment.

Number of Guesses: The total number of card pairs flipped before the end of the task. The number of card pair guesses indicated how frequently the player was mismatching pairs and how often they needed to revisit a previous card to recall its value and suit.

Errors: The total count of card pairs flipped that included only cards previously shown before completing the task. This count does not include card pairs that were correctly matched. Errors acted as an additional performance measure alongside the count of card pair guesses. Nevertheless, it reduced the element of chance present in card pair guesses. The number of card pair guesses could be fewer if the participant gets consecutive matches purely by chance. On the other hand, errors remained unchanged as they pertain to the participant recalling information they should already know.

Trial Length: The total time in seconds from flipping the first card until the completion of the task, when all pairs are found.

We employed a within-subjects, repeated measures design in which every participant was exposed to each experimental condition. The conditions were counterbalanced by random permutation, with a C# list shuffling algorithm.

3.4 Procedure

Each participant was greeted in the experiment room and given a brief overview of the experiment. Subsequently, the participant signed an informed consent form and completed a general pre-experimental questionnaire that gathered demographic information, experience with AR, and eyeglass usage. Afterward, we thoroughly described the experimental task (a concentration card game). The participant practiced by completing a round of the task on a smaller 2x2 grid consisting of 4 cards, continuing until the experimenter was satisfied that the participant had grasped the task. Hereafter, with assistance from the experimenter, the participant put on the Microsoft HoloLens 2 OST AR display. The participant completed the HoloLens 2 calibration process using the device's

built-in calibration software. Participants were instructed to ask any clarifying questions needed during the training period and provided a sheet describing the task they were trying to complete.

The participants then performed the task while sitting on a chair. Participants were notified that information would be shown on the right side of the cards via the AR headset depending on the experimental condition. However, participants were not notified which condition would be present each trial, or what all of the conditions prior to beginning the trials. There were no restrictions on the duration of each condition. Each session was recorded in video to monitor session length and other dependent variables. Upon finishing an experimental condition, the participant would respond to questions in a semi-structured interview that took approximately two minutes. During the semi-structured interview, the participants were asked three questions: 1. What are your thoughts on your performance in the game and the reasons for it? 2. Did you create any strategies to finish the game?, and 3. Was the AR helpful in accomplishing the game? The experimenter documented the responses of the participants for subsequent analysis. The participant was then expressed gratitude and given compensation if applicable. The complete experimental session lasted approximately one hour.

3.5 Participants

A total of 19 participants, comprising 13 men and 6 women, were recruited from the local university community. The ages of the participants ranged from 20 to 57 years, with an average age of 26 years. There were no restrictions for individuals with glasses or lenses; 10 participants wore corrective eyewear. The Institutional Review Board of the local university approved the study protocol. Participants were given the option to receive an Amazon gift card $20 USD or participate without compensation. Three participants were excluded from the study. Technical errors led to incomplete trials for two participants. In addition, a participant was a significant outlier. Consequently, our research contained 16 of the 19 participants.

Chapter 4

Results & Analysis [3]

In this chapter, we evaluate and discuss the results of the performed experiment. We performed both a thematic analysis of participant interviews as well as a Quantitative Analysis of the participant performance data including their trial length, their number of guesses, and their number of errors made in the trials.

4.1 Quantative Analysis

To evaluate the impact of our experimental conditions on participant performance, we conducted a *repeated measures ANOVA* using the ez package in R [11]. For effect size, we report the generalized eta squared (η^2), with interpretations as follows: small ($\eta^2 = 0.01$), medium ($\eta^2 = 0.06$), and large ($\eta^2 = 0.14$) [10]. In cases where the ANOVA was significant, we used the *Scheffe test* for post hoc comparisons between conditions because it has strong protection against Type 1 error. The Scheffe test was performed using the DescTools package in R [19]. The significant results of the Sheffe test are given in Table 4.1. Throughout the analysis, the p-value thresholds were set as follows: $* :< 0.05, ** :< 0.01, *** < 0.001$.

Card Pairs Guessed: Fig. 4.1 shows the results of the task performance under different conditions considering the metric of guessed pair of cards. ANOVA analysis revealed a significant effect of the conditions (main effect) on the pairs of cards guessed ($F_{4,60} = 8.61, p < .001, \eta^2 = 0.29$, *large*.) The Scheffe test revealed a significant difference between the AR helpful condition ($M = 13.44, SD = 1.21$) and the other experimental conditions: Combined ($M = 17.94, SD = 2.59, p < 0.05$); Irrelevant ($M = 19.19, SD = 3.25, p < 0.001$); Misleading

[3] Sturgeon, M., Anspach, E., Ortega, F., Ray, I., Safayet Arefin, M.: Impact of Relevant Augmented Reality Information on Human Performance. In: Bebis, G., Patel, V., Gu, J., Panetta, J., Gingold, Y., Johnsen, K., Arefin, M.S., Dutta, S., Biswas, A. (eds.) Ad- vances in Visual Computing. pp. 185–198. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-77389-1_15

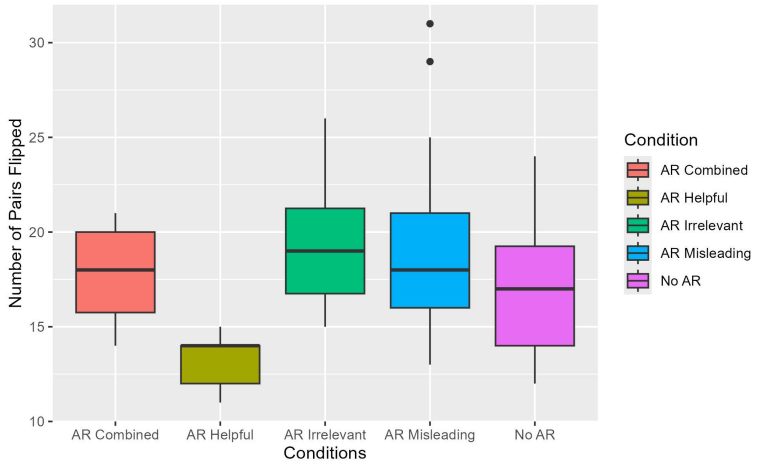

Figure 4.1: The average number of guesses/pairs made for each different AR condition type is shown. Participants in the AR Helpful condition made fewer guesses, indicating a positive impact on task performance compared to other conditions. Irrelevant and misleading AR information negatively effects the performance.

Table 4.1: Significant results of the Scheffe test for all dependent variables with lower and upper 95% confidence interval (CI).

Scheffe Test on Card Pairs Guessses					
Compared Conditions	diff	Lower CI	Upper CI	p value	Significance
Irrelevant-Helpful	5.750	1.90737102	9.592629	0.00055	***
Misleading-Helpful	5.875	2.03237102	9.717629	0.00039	***
Combined-Helpful	4.500	0.65737102	8.342629	0.01267	*
No AR-Helpful	3.875	0.03237102	7.717629	0.04698	*
Scheffe Test on Errors					
Irrelevant-Helpful	2.8750	-0.1265856	5.876586	0.0677	.
Misleading-Helpful	3.1250	0.1234144	6.126586	0.0367	*
Scheffe Test on Trial Lengths					
Irrelevant-Helpful	39.6250	1.438953	77.811047	0.0377	*
No AR-Irrelevant	-42.3125	-80.498547	-4.126453	0.0216	*

($M = 19.31, SD = 5.15, p < 0.001$); and No AR ($M = 17.31, SD = 3.75, p < 0.05$). No significant differences were found between the other conditions. This suggests that the presence of helpful AR content led participants to flip fewer cards during the task, thereby enhancing their

performance. Also, misleading and irrelevant AR information decreased the task performance in similar manner.

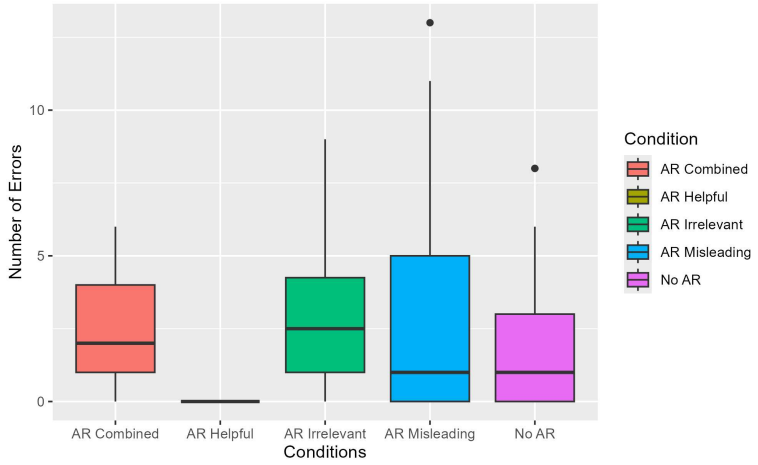

Figure 4.2: Average number of errors made in a round separated by each different AR condition type is presented. The AR Helpful condition has a smaller average and lacks distribution in comparison to the other conditions, indicating that irrelevant and misleading AR information detrimentally impacts performance.

4.1.1 Errors

The task performance effect of errors and conditions are reported in Fig. 4.2. According to ANOVA, there was a significant effect of conditions on the number of errors ($F_{4,60} = 3.76, p < .01, \eta^2 = 0.15, \textit{large}$). There was a significant difference between the AR helpful condition and AR misleading condition ($M = 3.13, SD = 4.08, p < 0.05$), and closed to significance versus the irrelevant condition ($M = 2.88, SD = 2.7, p = 0.06$). No significant differences were found between the other conditions. Thus, the AR helpful condition led to a reduction in errors compared to the misleading and irrelevant conditions. This suggests that misleading and irrelevant AR information led to poorer task performance than helpful AR information and similar performance to when no AR overlay was provided.

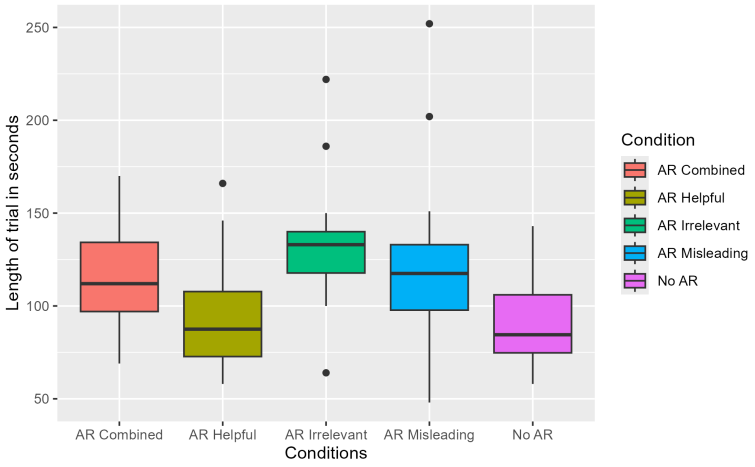

Figure 4.3: The figure illustrates the mean trial durations in seconds categorized by various AR condition types. The AR irrelevant condition exhibits the highest average trial time and the broadest range of times, followed by AR misleading , shown by the mean line on each boxplot, and the distribution of the data by the outlier points. Participants completed trials more faster under AR helpful conditions and with no AR overlay.

4.1.2 Trial Length

Fig. 4.3 shows the results of the trial length at each of the experiment conditions. Our ANOVA analysis revealed that the differing conditions had a largely significant effect on the trial length ($F_{4,60} = 7.48, p < .001, \eta^2 = 0.20, \textit{large}$). In the Scheffe test, the AR irrelevant condition ($M = 133.69, SD = 34.73$) was significantly worse than the non-AR condition ($M = 91.38, SD = 24.51, p < .05$). Moreover, AR helpful condition ($M = 94.06, SD = 29.79$) was significantly better than the irrelevant condition ($p < 0.05$). No significant differences were found between the other conditions. Thus, the AR irrelevant condition resulted in an increase in the trial length duration compared to the AR helpful condition and when no AR overlay was presented.

4.2 Thematic Analysis

For our qualitative analysis, we utilized thematic analysis based on the recorded documentation from the semi-structured interviews we conducted. We reviewed the interview data for each participant, extracted information, developed codes, and established the themes. Overall, two themes

emerged from our interview data: *strategies* and *views on AR conditions*. The narrative of each theme is described below.

Strategies:

Utilizing Hurtful AR: Several participants mentioned using irrelevant, misleading, or combined overlays to identify cards that had been seen before. This helped them keep better track of the whereabouts of unseen cards in the play area. Additionally, some participants effectively utilized the accurate overlays from the combined condition to enhance their performance.

Ignoring Hurtful AR: Several participants reported being more successful during the irrelevant, misleading, or combined overlays by disregarding the content shown to them. This observation varied between participants and conditions. Some participants quickly recognized that the overlay was unhelpful and ignored it with little effect on their performance. Others needed more time to identify and disregard the overlay, negatively affecting their task performance.

Learning Effect: Some participants observed that even with different overlays, their performance in the card game enhanced over the rounds. Repetition of the task enabled them to refine and enhance their strategy with each round.

One trend seen in some of the interviews was a belief that the AR content was meant to help them some even describing it as giving them "hints", even when it was clearly incorrect or unrelated to the current task. This suggests a possible underlying trust in devices providing helpful or useful information when being used, and that users may try to interpret the information as helpful if possible.

Views on AR Conditions:

AR Helpful: Most participants felt they did well with the helpful overlays because they continuously showed the correct card information once it had been flipped. Three participants

even likened this overlay to cheating, as it made the game too easy. Some participants also mentioned that they started to depend on the helpful AR information once they trusted its consistency.

AR Irrelevant: Many participants reported difficulties playing the game because irrelevant information distracted them from their objective. Some even mentioned that they thought there was a link between pairing irrelevant facts and cards, which caused them to make errors and forget previously seen information.

AR Misleading: A number of participants indicated that their performance declined because the misleading information led them to confuse the cards they had memorized. Conversely, other participants found this overlay to be unproblematic as they quickly disregarded the displayed information once they identified it as false.

AR Combined: Some reported a negative impact on their performance, with one person noting that recalling overlays with accurate information strained their memory. Others felt indifferent, as they could easily disregard the mixed information. Conversely, some found it beneficial because they could reliably trust parts of the overlay's information.

No AR: Several participants mentioned that the lack of overlays made it simpler to focus on the game. Despite this, some participants made errors in this condition; one noted that they had become too dependent on the AR overlays and had forgotten how to play the game without them.

4.3 Discussion

Our first hypothesis (H1) stated that helpful AR information would improve task performance. Our statistical analysis shows that the helpful AR condition helped users complete the task more effectively compared to other conditions, with participants making fewer guesses in the helpful condition. When examining errors, there were more errors among participants when AR information was absent compared to the helpful AR condition, although the difference was not significant.

Interestingly, the average duration of the trial length was shorter without the AR overlay than with the helpful AR information. One possible explanation is that participants spent additional time reading the AR overlay, although it was beneficial. Consequently, our findings partly support hypothesis H1. Furthermore, our thematic analysis indicated that most of the participants were satisfied with their performance under the helpful AR condition, suggesting that the AR information directly assisted the participants in recalling the positions of the cards by providing them with answers.

Our second hypothesis (H2) stated that misleading information would negatively affect human performance. Our findings confirm H2. We observed that the misleading conditions reduced participant performance, suggesting that the manipulative content caused delays and impaired participants' ability to quickly recall card positions. A potential reason is that participants took extra time to distinguish their memory of the card from the information shown on the overlay under the misleading condition. For irrelevant information, participants may have been slowed down by spending additional time reading the AR overlays, which was not related to the task.

Our third hypothesis (H3) stated that irrelevant information would negatively affect human performance. Our findings confirm H3. We observed that the irrelevant condition reduced participant performance, suggesting that the manipulative content introduced delays or distracted the participants, thus impairing the participant's ability to recall card positions in a timely manner. A potential reason for this is visual obstruction on the card from the overlay, or extra time spent processing unrelated information similar to a context switch. Participants mentioned that during the irrelevant condition they began to disregard the overlays once they realized that the information was not helpful. The additional time required to reach this understanding might have further delayed the completion of their task.

Our fourth hypothesis (H4) proposed that participant performance would be comparable when provided with a combination of misleading, irrelevant, and useful AR information with performance observed when only misleading or irrelevant information is presented. The results we obtained confirm H4. No significant differences were found between the combined AR condition

and the conditions with only misleading or irrelevant AR information. Our finding implies that the user's memory performance is similarly affected by the presence of misleading and irrelevant content along with helpful AR information. Participants were similarly affected when exposed to partially manipulative content compared to fully manipulative content.

Comparison to prior work These results align with what we would expect based on the previous literature, Cheng et al. [5] found that perceptual attacks have a negative impact on human metrics such as reaction time, situational awareness and sustained attention, all of which are important features of human perception. Of particular note to our work is their result on sustained attention, where they found that auditory perceptual manipulation attacks may have resulted in additional failures over participants who did not experience the manipulative content in memorizing a sequence of real world stimuli. For our work this reflects the result that we found in our manipulative conditions where the ability to memorize and recall card positions was negatively impacted by the presence of manipulative content, however where they focused on an auditory distraction our work focused on how the quality of information could be exploited. They also reported similar responses around belief in the usefulness of devices that we found, where participants described the motivation behind the XR content as helpful even when the information provided itself was not helpful.

Our work also aligns with the results found in both Tseng et al. and Casey et al. [4, 22]. They outlined attacks that may influence changes in user behavior such as changes in movement or other decisions. Our work aligns with the effects that they noted with other similar attacks such as Casey's Overlay attack. They found that malicious overlays could distract or impair user actions within a virtual environment based on what kind of content was presented. They did not focus on the specific granularity of information quality but rather on categories of information that could cause harm such as excessive gore or sexual content; however it is reasonable to believe a similar result is likely to be seen at the granular level within a single category as we found in our results.

4.3.1 Attack Success

An attacker performing a Perception Manipulation Attack on the device would have likely been successful at degrading and obstructing user performance in this case. An attacker could utilize a similar setup to maliciously change the quality of the information presented to a AR user to slow them down or introduce errors in the system. Take for instance AR surgery where the information presented to the user is paramount to correct treatment and the task is often very time sensitive. The introduction of errors or lengthening of task times as seen in our work could have disastrous consequences to the patient.

Chapter 5

Conclusion and Future Work [4]

We examined how manipulative information in OST AR affects and influences human performance. We carried out an experiment in which the participant conducted a card-concentration game as a task. The participant engaged with various kinds of AR information, which could be helpful, misleading, irrelevant, or a mixture of these three. We conducted interviews to understand user views and approaches to dealing with this information to draw inferences as to what kind of attacks would be effective against users. The primary findings of our investigation are: manipulative AR information degraded task performance and helpful AR information increased task performance in the OST AR system. It means that manipulative content, such as irrelevant or misleading, impacts a user's ability to complete a task by breaking their attention from the main task, causing them to take active steps to address and avoid manipulative content. Our work provides insights into how a Perceptual Attack that utilizes information quality may be able to degrade user performance within the system.

5.1 Future Work

Our study has multiple limitations which offer directions for future research, as well as our results suggest directions for additional study.

Learning Effect: We detected the presence of a learning effect. As each participant went through the conditions sequentially, the latter conditions were inclined to yield better results as the participants honed their strategies. Consequently, future research might consider a between-subject design with a larger sample size. This would allow the effect of each individual condition to be

[4] Sturgeon, M., Anspach, E., Ortega, F., Ray, I., Safayet Arefin, M.: Impact of Relevant Augmented Reality Information on Human Performance. In: Bebis, G., Patel, V., Gu, J., Panetta, J., Gingold, Y., Johnsen, K., Arefin, M.S., Dutta, S., Biswas, A. (eds.) Ad- vances in Visual Computing. pp. 185–198. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-77389-1_15

surveyed separately of eacho ther preventing participant learning in the system between trials and trial conditions.

Device Trust: Additionally, the interaction between the user trust in the device and the effectiveness of the device needs to be investigated further. Participants with a greater trust in the device may be effected to a greater magnitude when manipulative content is introduced due to more closely following the virtual information. A future study could provide a more complex task where the user is made familiar and more reliant on the AR technology to assist them prior to the introduction of manipulative content. This would be more reflective of common cybersecurity scenario's where a user commonly uses a device and relies on its functionality, and would provide a more accurate picture how effective cyberattacks of this nature are at disrupting XR device usefulness,

Tracking Focus Time: In our experiment, we did not use any eye-tracker and were unable to determine how much time participants were focusing on different types of AR information. It remains unclear whether participants focus more on the AR content when the information is manipulative. Therefore, future studies should incorporate eye-tracking to address this issue. Additionally, this would offer further context as to why user performance was affected and could verify feedback from interviews suggesting that the AR content was obstructive.

Number of Manipulative Contents: In our study, the participant had to handle just a single piece of AR information for each card. It is essential to determine the volume of manipulative content required to affect human performance. If minor manipulations are employed in an attack, it might be more difficult for users to precisely discern whether an attack is happening or if it's merely a technical bug or glitch.

Identifying Malicious Distraction Notifications and overlays in XR environments can be used to introduce content to a user, this content could be used by malicious actors and attackers to

distract a user or to introduce harmful content. Future works should study systems to identify when malicious or distracting content is introduced to a user to provide mitigation or protection strategies.

Simulated Cyberattack To further identify the impact of a cyberattack compromising information quality a future study must be conducted where the participants are placed in a real scenario and manipulative content is introduced by a simulated attacker similar during the course of the task as a live attack scenario. This can be performed with a variety of tasks such as simulated VR surgery, AR plumbing or military simulation to investigate possible consequences of and user reactions to such an attack.

Cyberattacks against device experts A field expert using a device may have a different relationship with the task performed and the XR device used as an aid on that task. This may cause information quality manipulations to have a greater impact due to them being able to rely on their own personal expertise or lesser impact due to their comfort and experience using the device. Future work should investigate how these attacks and degradation of information quality as a whole impacts more experienced expert participants. This would provide a more accurate view of how fields such as the military, surgery or engineering; where the members are experienced professionals would be impacted by these attacks.

References

[1] Arafat, A.A., Guo, Z., Awad, A.: VR-Spy: A Side-Channel Attack on Virtual Key-Logging in VR Headsets. In: 2021 IEEE Virtual Reality and 3D User Interfaces (VR). pp. 564–572 (Mar 2021). https://doi.org/10.1109/VR50410.2021.00081, https://ieeexplore.ieee.org/abstract/document/9417659, iSSN: 2642-5254

[2] Azuma, R.T.: A survey of augmented reality. Presence: Teleoper. Virtual Environ. **6**(4), 355–385 (Aug 1997). https://doi.org/10.1162/pres.1997.6.4.355, http://dx.doi.org/10.1162/pres.1997.6.4.355

[3] Billinghurst, M., Duenser, A.: Augmented reality in the classroom. Computer **45**(7), 56–63 (July 2012). https://doi.org/10.1109/MC.2012.111

[4] Casey, P., Baggili, I., Yarramreddy, A.: Immersive Virtual Reality Attacks and the Human Joystick. IEEE Transactions on Dependable and Secure Computing **18**(2), 550–562 (Mar 2021). https://doi.org/10.1109/TDSC.2019.2907942, https://doi.org/10.1109/TDSC.2019.2907942

[5] Cheng, K., Tian, J.F., Kohno, T., Roesner, F.: Exploring User Reactions and Mental Models Towards Perceptual Manipulation Attacks in Mixed Reality. Proceedings of the 32nd USENIX Security Symposium (2023)

[6] Doshi, A., Smith, R., Thomas, B., Bouras, C.: Use of projector based augmented reality to improve manual spot-welding precision and accuracy for automotive manufacturing. The International Journal of Advanced Manufacturing Technology **89**, 1279–1293 (03 2017). https://doi.org/10.1007/s00170-016-9164-5

[7] Gong, Z., Wang, R., Xia, G.: Augmented Reality (AR) as a Tool for Engaging Museum Experience: A Case Study on Chinese Art Pieces. Digital **2**(1), 33–45 (Mar 2022).

https://doi.org/10.3390/digital2010002, https://www.mdpi.com/2673-6470/2/1/2, number: 1 Publisher: Multidisciplinary Digital Publishing Institute

[8] Gulhane, A., Vyas, A., Mitra, R., Oruche, R., Hoefer, G., Valluripally, S., Calyam, P., Hoque, K.A.: Security, Privacy and Safety Risk Assessment for Virtual Reality Learning Environment Applications. In: 2019 16th IEEE Annual Consumer Communications & Networking Conference (CCNC). pp. 1–9 (Jan 2019). https://doi.org/10.1109/CCNC.2019.8651847, https://ieeexplore.ieee.org/document/8651847, iSSN: 2331-9860

[9] Gürlük, H.: Concept of an adaptive augmented vision based assistance system for air traffic control towers. In: 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC). pp. 1–10 (Sep 2016). https://doi.org/10.1109/DASC.2016.7777975, https://ieeexplore.ieee.org/document/7777975/?arnumber=7777975, iSSN: 2155-7209

[10] Lakens, D.: Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. Front Psychol **4**, 863 (Nov 2013), https://pubmed.ncbi.nlm.nih.gov/24324449/

[11] Lawrence, M.A.: ez: Easy Analysis and Visualization of Factorial Experiments (2016), https://CRAN.R-project.org/package=ez, r package version 4.4-0

[12] Ling, Z., Li, Z., Chen, C., Luo, J., Yu, W., Fu, X.: I Know What You Enter on Gear VR. In: 2019 IEEE Conference on Communications and Network Security (CNS). pp. 241–249 (Jun 2019). https://doi.org/10.1109/CNS.2019.8802674, https://ieeexplore.ieee.org/document/8802674

[13] Livingston, M.A., Rosenblum, L.J., Brown, D.G., Schmidt, G.S., Julier, S.J., Baillot, Y., Swan, J.E., Ai, Z., Maassel, P.: Military Applications of Augmented Reality. In: Furht, B. (ed.) Handbook of Augmented Reality, pp. 671–706. Springer, New York, NY (2011). https://doi.org/10.1007/978-1-4614-0064-6_31, https://doi.org/10.1007/978-1-4614-0064-6_31

[14] Narzt, W., Pomberger, G., Ferscha, A., Kolb, D., Müller, R., Wieghardt, J., Hörtner, H., Lindinger, C.: Augmented reality navigation systems. Universal Access in the Information Society **4**, 177–187 (03 2006). https://doi.org/10.1007/s10209-005-0017-5

[15] Punpongsanon, P., Iwai, D., Sato, K.: SoftAR: Visually Manipulating Haptic Softness Perception in Spatial Augmented Reality. IEEE Transactions on Visualization and Computer Graphics **21**(11), 1279–1288 (Nov 2015). https://doi.org/10.1109/TVCG.2015.2459792, https://ieeexplore.ieee.org/document/7165660, conference Name: IEEE Transactions on Visualization and Computer Graphics

[16] Punpongsanon, P., Iwai, D., Sato, K.: FleXeen: Visually Manipulating Perceived Fabric Bending Stiffness in Spatial Augmented Reality. IEEE Transactions on Visualization and Computer Graphics **26**(2), 1433–1439 (Feb 2020). https://doi.org/10.1109/TVCG.2018.2871044, https://ieeexplore.ieee.org/document/8468063, conference Name: IEEE Transactions on Visualization and Computer Graphics

[17] Shi, C., Xu, X., Zhang, T., Walker, P., Wu, Y., Liu, J., Saxena, N., Chen, Y., Yu, J.: Face-Mic: inferring live speech and speaker identity via subtle facial dynamics captured by AR/VR motion sensors. In: Proceedings of the 27th Annual International Conference on Mobile Computing and Networking. pp. 478–490. MobiCom '21, Association for Computing Machinery, New York, NY, USA (Oct 2021). https://doi.org/10.1145/3447993.3483272, https://dl.acm.org/doi/10.1145/3447993.3483272

[18] Shuhaiber, J.H.: Augmented Reality in Surgery. Archives of Surgery **139**(2), 170–174 (Feb 2004). https://doi.org/10.1001/archsurg.139.2.170, https://doi.org/10.1001/archsurg.139.2.170

[19] Signorell, A.: DescTools: Tools for Descriptive Statistics (2024), https://CRAN.R-project.org/package=DescTools, r package version 0.99.54

[20] Sturgeon, M., Anspach, E., Ortega, F., Ray, I., Safayet Arefin, M.: Impact of Relevant Augmented Reality Information on Human Performance. In: Bebis, G., Patel, V., Gu, J., Panetta, J., Gingold, Y., Johnsen, K., Arefin, M.S., Dutta, S., Biswas, A. (eds.) Advances in Visual Computing. pp. 185–198. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-77389-1_15

[21] Sun, Q., Patney, A., Wei, L.Y., Shapira, O., Lu, J., Asente, P., Zhu, S., Mcguire, M., Luebke, D., Kaufman, A.: Towards virtual reality infinite walking: dynamic saccadic redirection. ACM Transactions on Graphics **37**(4), 67:1–67:13 (Jul 2018). https://doi.org/10.1145/3197517.3201294, https://dl.acm.org/doi/10.1145/3197517.3201294

[22] Tseng, W.J., Bonnail, E., McGill, M., Khamis, M., Lecolinet, E., Huron, S., Gugenheimer, J.: The Dark Side of Perceptual Manipulations in Virtual Reality. In: CHI Conference on Human Factors in Computing Systems. pp. 1–15 (Apr 2022). https://doi.org/10.1145/3491102.3517728, http://arxiv.org/abs/2202.13200, arXiv:2202.13200 [cs]