

# Exploring Head Tracked Head Mounted Displays for First Person Robot Teleoperation

Corey Pittman

University of Central Florida  
Orlando, FL 32816 USA  
cpittman@knights.ucf.edu

Joseph J. LaViola Jr.

University of Central Florida  
Orlando, FL 32816 USA  
jjl@eecs.ucf.edu

## ABSTRACT

We explore the capabilities of head tracking combined with head mounted displays (HMD) as an input modality for robot navigation. We use a Parrot AR Drone to test five techniques which include metaphors for plane-like banking control, car-like turning control and virtual reality-inspired translation and rotation schemes which we compare with a more traditional game controller interface. We conducted a user study to observe the effectiveness of each of the interfaces we developed in navigating through a number of archways in an indoor course. We examine a number of qualitative and quantitative metrics to determine performance and preference among each metaphor. Our results show an appreciation for head rotation based controls over other head gesture techniques, with the classic controller being preferred overall. We discuss possible shortcomings with head tracked HMDs as a primary input method as well as propose improved metaphors that alleviate some of these drawbacks.

## Author Keywords

3D Interaction; User Studies; Robots

## ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI):  
Miscellaneous

## General Terms

Design, Experimentation

## INTRODUCTION

As the cost of Virtual Reality technologies fall, new applications for these once unaffordable technologies are being found in a number of diverse areas of study. One such area is Human Robot Interaction. Manipulating robots using gestural inputs has been an oft studied area [3][4][10] in recent years, with interactivity and naturalness being among the observed benefits of this modality. A mode of control that has not been thoroughly explored as a primary input method for teleoperation is head tracking. Head tracked HMDs have often been used in virtual environments to increase a user's

sense of immersion and proprioception [3]. Using head tracking in addition to a head mounted display (HMD) as a means to improve a user's experience in virtual environments has been detailed in prior work, with evaluations of different settings to determine the effects of different HMD configurations being among the more recent work in the area [7][11].

One area that has not seen a significant amount of work with head tracking and head mounted displays is robot teleoperation. In this paper, we focus on direct teleoperation of a robot using egocentric head movements. The robot selected as the test robot for this study was a low cost quadrotor, as it is well represented in entertainment as a toy and possesses similarities to military unmanned aerial vehicles (UAV) in potential applications. Using commonly observed vehicles and VR control techniques as inspiration, we developed five head tracking based metaphors for flying the UAV using simply an HMD with head tracking and compared against a game controller. We designed and conducted a user study that asked users to complete a navigation task by flying around a hexagonal course. We then analyzed the qualitative and quantitative results of the study.

## RELATED WORK

Pfeil et al. [10] have previously developed full body gesture interfaces for controlling an AR Drone using a depth camera and joint information. Multiple metaphors for controlling a UAV were proposed, and a number of interfaces were compared in a formal user study. One significant problem with the system as described was the user's view of the drone, which was from a third person perspective. Our system addresses this problem by placing users in the first person, making our controls more meaningful to users given the difference in relative viewpoints.

Higuchi et al. [4] have developed a head tracking solution for the teleoperation of a Parrot AR Drone using synchronized optical trackers attached to the head and drone chassis called Flying Head. Users wore a head mounted display which featured a view from the front mounted camera of the drone, giving the user a first person view of the world through the drone's eyes. One of the techniques we implemented was based on this work with some modifications. However, we do not use any visual sensors thereby avoiding potential occlusion problems. The authors also did not conduct a systematic study comparing their approach with a traditional joystick interface.

Koeda et al. [6] tested the benefit of providing a user with a first person view from a teleoperated small helicopter. Their

system placed an omni-directional camera on the helicopter so that the user can turn his head while wearing an HMD and see as if looking around from within the helicopter. They did not utilize head tracking for any sort of teleoperation controls instead opting for an RC controller.

The Robonaut, a humanoid designed for use in space, was designed with a teleoperation technique for control of the entire robot, including the head with two degrees of freedom [1]. Users wore an HMD with head tracking and were able to send commands to the Robonaut by rotating head yaw and pitch. This control scheme is comparatively simple, as it only tracks two axes of rotation of the head, while our techniques makes use of three axes of head rotation and three translation axes. Other work has explored using head mounted displays for controlling robots, such as [8][9].

### **INTERACTION TECHNIQUES**

We developed five gestural navigation techniques, with one of them being inspired by prior work [4]. All techniques were created under the assumption that the user is standing. For all techniques that use head gestures, additional controls for takeoff, landing, and resetting the drone were placed on a Nintendo Wiimote to give a supervisor control of the drone in case of emergencies.

#### **Head Translation**

This technique was inspired by commonly used techniques for navigating a virtual environment by moving the upper body in the desired direction of motion [2]. In our context, this means that by taking a step forwards and shifting the head over that foot, the head will translate from its original position and the UAV will then move forwards. This can be done in all four cardinal directions, meaning the user needs only to shift one foot in a direction and then shift their weight over that foot to translate the UAV in that direction relative to its current heading. Once the user returns their foot and head to their original position, the UAV will cease moving and return to its resting position. The user can also move in the ordinal directions to cause the UAV to move along a combination of two cardinal directions. To turn the UAV, the user can turn their head to the left or right and the drone will turn until the user returns to looking forwards relative to their initial orientation. The user can stand on their toes to increase the elevation of the UAV and squat down slightly to decrease the elevation of the UAV.

#### **Head Rotation**

This technique was inspired by observing people playing video games in a casual environment. Occasionally when people play first person perspective video games, they tilt or turn their head slightly while playing in an attempt to will their avatar to move more than what the constraints of the game controls will allow [12]. To control the drone using these movements, we made use of the rotational axes of the head. If the user tilts his head forwards, the UAV will move forwards; tilts his head backward then the UAV will move back; and tilts his head to the left or right to cause the UAV to strafe in the corresponding the direction. The user can also turn his head to cause the drone to turn in the corresponding

direction. In order to move the drone up or down, the user stands on their toes or squats down slightly.

#### **Modified Flying Head**

Higuchi et al. designed an interaction technique based on a visual tracking system that allowed for a UAV to synchronize its movements with a user's movements [4]. A major shortcoming with this method of interaction is that the size of the user's control area would have to match the UAV's environment. A number of methods were proposed to alleviate this shortcoming. We made use of the turning technique proposed in their work, which uses direct manipulation of the UAV's heading based on the heading of the user, allowing them to turn 360 degrees within their control space. To cause the drone to turn, the user rotates his entire body the same amount that they would like the drone to rotate, which will rotate to match the movements of the user. This is the largest difference between this technique and the previous two described techniques. To move the drone forwards, the user needed to step forwards along the direction their head was facing. This remains true regardless of how turned the user is from their initial heading: shifting the head forwards will always move the drone forwards. The user can also shift their head in the other three major directions to translate the drone in other directions. To control elevation, the user can stand on their toes or squat down.

#### **Flying Car and Plane**

One initial assumption that was made when designing these techniques was that when a user is standing in a neutral position, the UAV should not be moving and therefore require the user to make a pronounced movement to trigger any sort of change in its behavior. In some cases, this could be considered a hindrance, including the control of vehicles that spend a majority of their time in motion, such as a car. Therefore we designed the Flying Car metaphor based around the idea that when sitting in a car that is engaged in a forward gear, it will move forwards slowly using the motion generated by its idling engine. When the user is standing idly, the UAV will move forwards at a slow pace. If the user shifts back, it is synonymous with applying the brakes in a car, therefore causing the UAV to come to a stop. If the user shifts forwards or presses the accelerator, the drone will accelerate at full speed. To turn the car, the user simply looks to the left or right with their head. The user may also look up to move the drone upward or look down to move the drone downwards—this being similar to theoretical controls for a flying car from popular entertainment media. Because a car cannot strafe, there is no ability to strafe in this control scheme. The Plane technique is similar to flying car except that instead of turning the head to control the drone's rotation, the user must tilt their head.

#### **Wiimote**

A classic control scheme was implemented as a baseline for comparison with our head tracking interfaces. Pfeil et al. showed that gestural interfaces were preferred to the packaged smartphone application for the AR Drone with regards to naturalness, fun, and overall preference [10]. For that reason we developed our own simple control configuration using

a Nintendo Wiimote. The D-pad was mapped to drone translations within the horizontal axis, elevation was mapped to the A and B buttons, and turning the drone was mapped to the 1 and 2 buttons. The other three commonly used controls (takeoff, land, and reset/emergency) were mapped to the remaining buttons on the controller. Figure 1 shows these settings mapped to a controller.




Figure 1. Control configuration for Wiimote interface.

## USER STUDY

We designed a user study around a simple navigation task for a robot in a physical environment. Our goal was to analyze the performance and ergonomics of our developed interaction techniques compared to a traditional interface.

## Subjects

Eighteen students (16 male, 2 female) from the University of Central Florida were recruited to participate in the study. Ages ranged from 18 to 38 with a median age of 22. Of all participants, 16 had experience with remote controlled vehicles, five had worn a head mounted display prior to participating in the study, and five had experience with position tracking devices. Two of the eighteen students were graduate students.

## Devices and Apparatus

We used an Oculus Rift, a low cost HMD with a stereoscopic display and a resolution of 1280 by 800. We used a Polhemus PATRIOT tracker to track the users head movements. The sensor was placed on the upper headband of the Oculus Rift and the source was placed at a position over the user's head to minimize electromagnetic interference from the electronics within the HMD. A Nintendo Wiimote was used for all non-gesture commands, in addition to being used as the basis for the control interaction technique for the study. The UAV we used to test our metaphors was the Parrot AR Drone 2.0. The AR Drone possesses two on board cameras: one forward facing and one downwards facing. The raw front camera feed from the UAV was displayed in the HMD with no additional modification or overlay. The system was set up on a laptop running Ubuntu 12.04 LTS, 4 GB of RAM, a 2.53 GHz Intel Core i5 processor, and an NVIDIA GeForce 310m graphics card.

## Environment Layout

The study took place in a 10m by 10m open area with six archways of varying heights arranged in a regular hexagon




Figure 2. The devices used for our system. From left to right: Laptop, Oculus Rift with Polhemus PATRIOT sensor attached, Parrot AR Drone 2.0, Nintendo Wiimote.

with a 3 meter gap between each arch. The archways were 1.2 meters wide and between 1.5 meters and 2 meters in height. These arches were arranged in alternating height along the hexagon. The size of these arches was determined in a pilot study in which circular hoops of various heights were placed in a similar environment. It was determined that a larger rectangular shape was more fitting for a flying navigation task. The participant stood 3 meters away from the course in the same room. The tracking system, HMD, and laptop were all set up in this designated area, allowing the participant to move around a designated area to give commands to the drone unimpeded. The participant's space was approximately 1.5 meters by 1.5 meters, giving ample space to freely execute head gestures without impedence. Though they could hear the drone in the environment, participants were asked to look away from the course and towards the control setup at all times. The starting position of the drone was in an adjacent corner to the participant's position, approximately 6 meters away. A diagram illustrating this layout can be seen in Figure 3.




Figure 3. A diagram of the study environment layout.

## Experimental Procedure

We used a six-way within-subjects comparison where the independent variable was the interface being used to complete the course. The dependent variables were the completion time of the navigation task and responses to a post-questionnaire that rated a users agreement to nine statements on a seven point Likert scale for each of the six interfaces. These statements are listed in Table 1. A Simulator Sickness Questionnaire (SSQ) [5] was also completed after each run to check for any sort of nausea caused by the use of a head mounted

display. The six interfaces tested were Head Translation, Head Rotation, Modified Flying Head, Flying Car, Plane, and Wiimote. The interface order was assigned using replicated Latin Squares.

Participants were given a practice run at the course to become familiar with the nuances of the current control scheme. Once the participant declared that they felt confident in their abilities to complete the course for time, the drone was placed back at the initial point of the course. This typically occurred after one run, though there were examples of participants needing upwards of three practice runs to feel comfortable enough with the controls to be willing to attempt a timed run. With the drone at the starting point of the course, the proctor of the study used the Wiimote to give the takeoff command to the drone. Once the drone was in the air, the participant was informed that time would start on their mark. Participants were allowed to make heading adjustments to the drone prior to the start of a run to correct for any unintended movement or drift that occurred on takeoff.

Once the participant gave the command to move the drone forwards, the proctor began timing the run. The participant was asked to move the drone through the six arches in counter clockwise order. If at any time the participant passed by an arch without flying through it, they were asked to continue on to the next arch and proceed through an additional arch after passing through what would have been the sixth arch. If the participant crashed the drone mid-run, there was no time penalty: the timer was stopped and the drone was returned to a point on the optimal flight path. Timing was continued once the participant began moving the drone forwards again. No penalty was given to participants who touched the arches while passing through them without crashing. Time was stopped once the entire drone passed through the final arch. A human was placed at the position of the last arch to ensure that an accurate time was recorded. Participants were given a break after each run to fill out a questionnaire.

When the Wiimote control interface was used to complete a run, the participant did not wear the HMD. They instead were asked to sit in front of a monitor with the controller in their hands. This was done to ensure that the interface we were comparing to most accurately represented a standard interface, which typically allow participants to freely look at the controller in their hands to confirm the button configuration.

## ANALYSIS OF DATA

### Quantitative Metrics

Figure 4 shows the mean completion time of the course for each technique. Using a 6-way repeated measures ANOVA analysis, we tested for significant differences in the mean completion times between each of the six interaction techniques. We found that there was significance ( $F_{5,13} = 5.885, p < 0.05$ ), and therefore went on to compare the Wiimote control technique to each of the head gesture based techniques. The Wiimote was found to be significantly faster around the course than each of the gesture based techniques based on a paired samples t-test of the mean completion times with  $\alpha = 0.05$ .

Between Runs Survey Questions
The interface to fly the Drone was comfortable to use.
The interface to fly the Drone did not confuse me.
I liked this interface for flying the Drone.
The interface to fly the drone felt natural to me.
It was fun to use this interface to control the drone.
I did not feel frustrated using the interface.
I did not feel discomfort or pain while using this interface.
It was easy to use this interface.
The Drone always flew in a predictable way.

Table 1. Survey questions asked after each run of the study. The questions used a 7-point Likert scale, with one being strongly disagree and seven being strongly agree.

To determine if there was any significant difference among the remaining techniques, we then ran a paired samples t-test with each of the remaining pairings. We found that there was significant difference between Plane and Flying Head ( $t_{17} = 2.858, p < 0.05$ ), Plane and Head Translation ( $t_{17} = 2.353, p < 0.05$ ), Head Rotation and Head Translation ( $t_{17} = 2.617, p < 0.05$ ), and Head Rotation and Flying Head ( $t_{17} = -2.457, p < 0.05$ ). This implies that though they were not nearly as fast as the Wiimote, Head Rotation and Plane were both better than Flying Head and Head Translation.

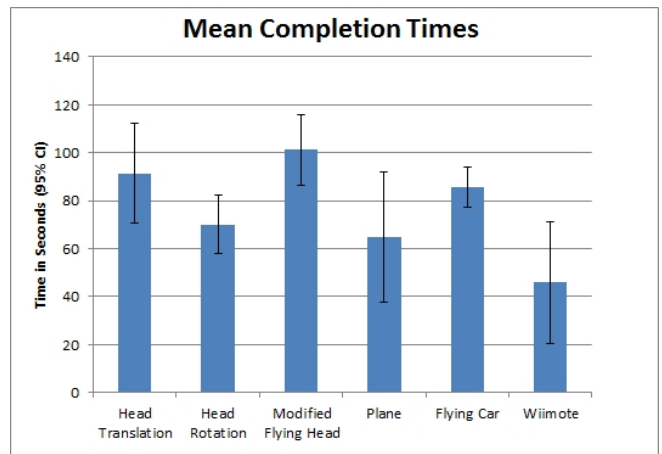


Figure 4. Mean completion time for each technique. There are three tiers of performance displayed here with Wiimote being significantly faster than all others, and Head Rotation/Plane being faster than the remaining three.

### Qualitative Metrics

To determine if there was any significant difference between median results on our qualitative ranking metrics, we used a non-parametric Friedman test. If it was determined that there was significant differences, we then employed Wilcoxon's signed rank test to look at significant differences between the Wiimote technique and the head tracking based techniques. We found that the Wiimote was rated highest along a number of different measures.

Wiimote was significantly more predictable, easy to use, and comfortable than every other interface. These findings match up with the completion time results, as navigation tasks will

typically be best completed with an interface that is accurate and familiar. Participants had a clear preference for the strain-free button presses of the controller overhead movements. Wiimote was significantly less frustrating, less confusing, more fun, and more likable than Modified Flying Head and Plane, based on analysis of the medians of ranking. One possible explanation for this is the large number of users who have previously used a remote control vehicle. All eighteen of the users also own at least one game system, so there is familiarity with similar controls to the one that was given to them in this experiment.

In order to determine a ranking of preference among participants, we used a Chi-Squared test to determine which techniques were liked more than expected. The Wiimote is ranked as the most likable interface of the six used in the study ( $\chi_{5,17}^2 = 21.333, p < 0.05$ ). The most liked of the remaining techniques is shown to be the Head Rotation interface ( $\chi_{5,17}^2 = 12.667, p < 0.05$ ). Head Translation and Modified Flying Head were the least liked interfaces.

Analysis of the SSQ scores shows significant difference among the six interfaces. The Wiimote was determined to have lower total SSQ scores than Head Rotation ( $t_{17} = -2.250, p < 0.05$ ), Plane ( $t_{17} = -3.296, p < 0.05$ ), and Flying Car ( $t_{17} = -2.148, p < 0.05$ ). There is significantly less sickness in the Wiimote technique than three of the head-tracked HMD interfaces, most likely because participants are not wearing an HMD when using it. Only five of the participants reported experiencing any amount of sickness that achieved a total SSQ score greater than 25 with any of the techniques. Eye strain scores were more prominent than any other measure, most likely due to the low resolution of the HMD and the motion blur from the camera feed. Another interesting pattern is that anytime a user had to wear the HMD for long periods of time to complete a run, their SSQ scores were notably higher. We base this observation solely on the completion time of the run, though participants typically wore the HMD for longer than the reported numbers when factoring in practice time.

A number of comments about the interfaces were recorded by participants between runs and after the study. Four of the 18 participants explicitly mentioned enjoying the natural feeling of the head tracked interfaces. Two each mentioned that they found the head gestures a fun and intuitive method of controlling a robot. Three participants mentioned that the drone felt like an extension of their body, verifying that they indeed felt immersed while wearing the HMD. There were two participants who explicitly mentioned that Head Rotation is most likely the most interesting of the interfaces, while simultaneously requiring the least energy to use.

Seven participants reported a lack of sensitivity or precision when using the head gestures. One explanation for this could be incorrect calibration, as participants tended to calibrate their head turning thresholds at an exaggerated position even after being informed that all thresholds should be comfortable and easily met using ordinary head movements. Only two of the participants reported having problems with the HMD being blurry in written form, though more than two mentioned

it while running the studies. Multiple participants mentioned drifting, difficulty returning to their original position while wearing the HMD, while using the Head Translation and Flying Head techniques. These problems stem from the use of an HMD, which completely occludes the users immediate surroundings, thereby preventing them from returning to their original position based on environmental references.

## DISCUSSION

Two interaction techniques (Flying Car, Plane) were based on metaphors that could not be meaningfully used in a large number of interaction contexts, as they assume that the robot will be moving forwards and therefore they are most suited for controlling robots designed around their namesakes: cars and planes. Our study did show some shortcomings when comparing these interfaces, particularly the lack of precision that participants mentioned. Our concerns about the size of the arches users were expected to navigate through with a limited field of view were confirmed to be true even with the larger archways. A larger FOV would have helped users to perform the task, as users often felt that they were smaller than the onboard camera which did not give a proper sense of the scale of the robot. These findings were expected, as prior studies have found that a larger field of view assists users in a number of different tasks [3] [11]. The other problem users had with the HMD was low resolution. Users stated that they had trouble discerning distant objects from their surroundings. When using the Wiimote, users did not experience any loss of visual information. This could have contributed to the users' preferences.

A few problems were caused by the AR Drone itself. Because we relied solely on the onboard magnetometer for heading information, interfaces that relied on it occasionally stuttered or wobbled anytime there was loss of information. The only time this occurred was when Flying Head was the interface being tested. There were oscillations in the sensor information, which caused a side to side vibration in the UAV as it turned while using Flying Head. Both of these contributed to the general disfavor of the Flying Head interface. The UAV also tended to drift to one side when given a forwards command with too small of a magnitude, which particularly affected Plane and Flying Car, as they were the only two techniques that primarily employed slower speeds.

We believed that Head Translation and Flying Head would be easy to use because stepping in a direction seemed to be a simple command for people to give. People easily drifted away from their starting position when they were unable to see their surroundings and then lose their original heading, therefore causing their controls to be less responsive. There was some praise for the one-to-one turning controls of the Flying Head technique, but combined with the recentering problem, the controls became inaccurate or unresponsive quickly. The large movements required to trigger a command also caused problems for users who wanted to give rapid commands to the UAV. A possible solution to these problems would be to combine the turning of these two techniques, allowing for direct control for fine adjustments of heading and a

point at which the drone will begin turning continuously until the user returns to a point below a threshold head yaw value.

Participants found Head Rotation to be a fun and responsive technique. A couple of participants found this interface to be more nausea inducing than other techniques, though not enough to negatively impact performance. This technique was originally designed to be used while sitting, and a number of users commented that it would have been easier to use while sitting. This mapping would be easiest to adapt to a number of different robot configurations and, if used from a sitting position and combined with a higher resolution HMD, could be used for extended periods of time

Both Flying Car and Plane received middling preference due to quirks in their design. Users felt that the lack of strafing and moving backwards was restricting and made it difficult to move the UAV in certain situations, such as when the user was stuck on an arch or wall. That Plane recorded such low completion times can be attributed to the automatic forwards movements of the interface. Users spent less time aiming for the arch and more time adjusting on the fly when they did not have to give a separate command to move forwards. This promoted a more reckless approach to the course that was also seen with Flying Car. Although these techniques would not be viable for many other interfaces, it is clear that when applicable, it is worth determining what sort of effect having automatic forwards movement would have on a user's experience.

Generally, our results lead us to believe that so long as the mapping is contextually sound, head tracking is worthy of exploring when attempting to teleoperate a flying robot. Though some findings of our study are limited to UAVs, we can infer that others will equally benefit other types of robots. The key design requirements for such an interface would be to limit the magnitude of the movements of the head to preserve user comfort. Ensuring that there is a logical correspondence between head movements and robot movements allows for a natural connection between the user and the teleoperated robot. This can be literal as in the direct axial mapping of Head Rotation or symbolic as in the Flying Car and Plane interfaces.

In the future, we aim to explore how to improve our head tracking techniques by improving their precision and adding stereoscopic view of the environment. We would also like to look at the possible design implications of adding an additional sensor to the user's torso. We also plan on making use of the Head Rotation technique to control the head of a humanoid robot along with a system for tracking the skeleton to allow for direct teleoperation of the humanoid in a human-robot team.

## CONCLUSION

We developed head tracking based interaction techniques using a Polhemus Patriot 6DOF tracker attached to an Oculus Rift HMD to control an AR Drone 2.0. Following a user study designed to evaluate the performance of the technique based on a number of quantitative and qualitative metrics, we found that users felt that a traditional game controller interface was

superior to our head tracking interfaces. Users consistently selected the Wiimote along almost all measures. However, users did find that Head Rotation was the most likeable of the non-traditional interfaces, although it did not perform as well as the Wiimote based on completion time and preferences. We can say that Head Rotation is a viable interface for controlling a quadrotor when users stay within the bounds of calibration. From this, we infer that Head Rotation is a possible alternative control scheme for a number of robot configurations, due to the many parallels that could be drawn between head rotations and movements of the robot platform.

## ACKNOWLEDGEMENTS

This work is supported in part by NSF CAREER award IIS-0845921 and NSF awards IIS-0856045 and CCF-1012056.

## REFERENCES

1. Bluethmann, W., Ambrose, R., Diftler, M., Askew, S., Huber, E., Goza, M., Rehnmark, F., Lovchik, C., and Magruder, D. Robonaut: A robot designed to work with humans in space. *Autonomous Robots* 14, 2-3 (2003), 179-197.
2. Bowman, D. A., Kruijff, E., LaViola Jr, J. J., and Poupyrev, I. *3D user interfaces: theory and practice*. Addison-Wesley, 2004.
3. de Vries, S. C., and Padmos, P. Steering a simulated unmanned aerial vehicle using a head-slaved camera and hmd. In *AeroSense '97*, International Society for Optics and Photonics (1997), 24-33.
4. Higuchi, K., and Rekimoto, J. Flying head: a head motion synchronization mechanism for unmanned aerial vehicle control. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, ACM (New York, NY, USA, 2013), 2029-2038.
5. Kennedy, R. S., Lane, N. E., Berbaum, K. S., and Lilienthal, M. G. Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology* 3, 3 (1993), 203-220.
6. Koeda, M., Matsumoto, Y., and Ogasawara, T. Development of an immersive teleoperating system for unmanned helicopter. In *Robot and Human Interactive Communication, 2002. Proceedings. 11th IEEE International Workshop on*, IEEE (2002), 47-52.
7. Lee, S., and Kim, G. J. Effects of haptic feedback, stereoscopy, and image resolution on performance and presence in remote navigation. *International Journal of Human-Computer Studies* 66, 10 (2008), 701-717.
8. Martins, H., and Ventura, R. Immersive 3-d teleoperation of a search and rescue robot using a head-mounted display. In *Emerging Technologies & Factory Automation, 2009. ETFA 2009. IEEE Conference on*, IEEE (2009), 1-8.
9. Mollet, N., and Chellali, R. Virtual and augmented reality with head-tracking for efficient teleoperation of groups of robots. In *Cyberworlds, 2008 International Conference on*, IEEE (2008), 102-108.
10. Pfeil, K., Koh, S. L., and LaViola, J. Exploring 3d gesture metaphors for interaction with unmanned aerial vehicles. In *Proceedings of the 2013 international conference on Intelligent user interfaces*, IUI '13, ACM (New York, NY, USA, 2013), 257-266.
11. Ragan, E., Kopper, R., Schuchardt, P., and Bowman, D. Studying the effects of stereo, head tracking, and field of regard on a small-scale spatial judgment task.
12. Wang, S., Xiong, X., Xu, Y., Wang, C., Zhang, W., Dai, X., and Zhang, D. Face-tracking as an augmented input in video games: enhancing presence, role-playing and control. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM (2006), 1097-1106.