How Confidence Information Influences Trust and Reliance in Human-Robot Teams

by

Alexandra Wolff

A Thesis Presented in Partial Fulfillment
of the Requirements for the Degree
Master of Science

Approved April 2022 by the
Graduate Supervisory Committee:

Nancy Cooke, Chair
Erin Chiou
Rob Gray

ARIZONA STATE UNIVERSITY

May 2022

ABSTRACT

Human-robot teams (HRTs) have seen more frequent use over the past few years, specifically, in the context of Search and Rescue (SAR) environments. Trust is an important factor in the success of HRTs. Both trust and reliance must be appropriately calibrated for the human operator to work faultlessly with a robot teammate. In highly complex and time restrictive environments, such as a search and rescue mission following a disaster, uncertainty information may be given by the robot in the form of confidence to help properly calibrate trust and reliance. This study seeks to examine the impact that confidence information may have on trust and how it may help calibrate reliance in complex HRTs. Trust and reliance data were gathered using a simulated SAR task environment for participants who then received confidence information from the robot for one of two missions. Results from this study indicated that trust was higher when participants received confidence information from the robot, however, no clear relationship between confidence and reliance were found. The findings from this study can be used to further improve human-robot teaming in search and rescue tasks.

ACKNOWLEDGEMENTS

I would like to recognize my academic advisor, Nancy Cooke, and my thesis committee, Erin Chiou and Rob Gray, for their knowledge, expertise, and patience throughout this journey. I would also like to recognize the HRI team of graduate and undergraduate researchers who conducted data collection, as well as the other CHART student researchers who have helped me during my time at Arizona State University.

TABLE OF CONTENTS

INTRODUCTION

Human-robot teams have been used to navigate dangerous and difficult environments in which humans alone cannot operate safely. For example, in an urban search and rescue environment, a human operator can work with a robot teammate to navigate rough and unsafe terrain. This requires team coordination, communication, and trust.

Human-robot teams have been used before in disaster environments. In June of 2021 following the Surfside condominium collapse in Miami, Florida, urban search and rescue responders utilized Unmanned Aerial Vehicles (UAVs) and semi-autonomous throwable military robots to initially evaluate and search the disaster site (Brown, 2021). This allowed the first responders to remain safe outside of the building while assessing the damage and evaluating the environment. The use of autonomous and semi-autonomous vehicles also helped with locating and ultimately recovering victims of the collapse.

In the future, human-robot teams will be highly interdependent and more advanced, thus it is essential that trust in the robot is properly calibrated. Future teams may also work with multiple robots, meaning that the operator must place trust in the robots and rely on them in order to successfully complete their missions. There are many theories and hypotheses on how to increase and properly calibrate trust in human-robot teams. One example is through the use of uncertainty or confidence information from the robot. However, the use of confidence level information in human-robot interactions is not well studied in complex task environments, such as in an urban search and rescue environment.

This thesis uses existing trust scales to look at the relationship between the presence of confidence level information and team trust in a simulated urban search and rescue task environment. The goal of this project is to further understand how confidence level information impacts trust in complex environments requiring non-binary decisions, and how future use could improve human-robot teams.

State of Knowledge: Trust and Reliance

Trust has been proven to be an essential component of human-robot teaming. According to Lee and See (2004), trust can be defined as “the attitude that an agent will help achieve an individual’s goals in a situation characterized by uncertainty and vulnerability.” Trust is important because it determines how a person may interact with an agent, and future usage of the agent (Billings et al., 2012). Previous research has shown that trust can be influenced by a multitude of factors. For example, Lewis and colleagues (2018) found that when workload is high, participants are more likely to overuse and over-trust an agent, increasing the likelihood of negative outcomes.

In human-robot teams, the human should trust the robot to perform a task if it is more suited to that task than themselves (Wang et al., 2018). However, if trust is not present or misplaced, disuse and misuse of the agent may occur. Disuse refers to the failures that occur when people reject automation’s capabilities when it would be correct and misuse refers to the failures that occur due to inappropriate reliance on automation, such as using it when they should not (Lee & See, 2004). Both are examples of inappropriate reliance on automation, which can have negative effects on performance and productivity. Trust in automation has been proven to affect reliance, as people tend to rely on automation they trust and not use automation they do not trust (Lewis et al., 2018). In order to have appropriate reliance and ultimately a high performing human-autonomy team, trust must be calibrated appropriately. In complex and uncertain environments, trust is likely to influence or guide reliance on an agent, especially when novel situations requiring adaptation occur (Lee & See, 2004).

Confidence Level Information

Transparency. Another factor that has been shown to impact trust in human-agent relationships is agent transparency. Chen and colleagues (2014) define agent transparency as “the descriptive quality of an interface pertaining to its abilities to afford an operator’s comprehension about an intelligent agent’s intent, performance, future plans, and reasoning process.” Based on this, they defined the Situation Awareness-based Transparency (SAT) model, which includes the three levels of transparency as they relate to the three levels of situation awareness. The first level of this model (SAT 1) consists of “the agent’s current state and goals, intentions, and plans. The second level (SAT 2) provides the agent's reasoning process and supports user comprehension of the agent’s actions. The third level (SAT 3) provides “information regarding [the agent’s] projection of future states, predicted consequences, likelihood of success/failure, and any uncertainty associated with the aforementioned projections” (Chen et al., 2016).

Transparency has been shown to improve operator trust in less reliable agents by revealing situations of high uncertainty, thus informing the operator of the agent’s limitations and developing appropriate reliance (Chen et al., 2016; Chen et al., 2014). Uncertainty information has also been shown to improve human-robot team performance (Chen et al., 2016).

For example, a study done by Wang and colleagues (2016) looked at the influence of transparency and explanations on a human-robot team task. In this study the difference between no explanations, explanation of some robot sensors, and confidence information were examined across human-robot teams of different abilities. The findings were that explanations that facilitated decision making improved trust, transparency, and performance in the robot, in conditions in which the robot's capability was limited. All explanation conditions, including confidence, were found to be useful in aiding the human teammate in deciding whether to trust the robot. The most interesting finding from this study, however, was that participants who received only confidence information for explanations felt as though they understood the robot's decision making process, despite the fact that they did not receive such information (Wang et al., 2016). This suggests that the use of confidence levels alone improves transparency, and that potentially combining confidence information with other explanations may significantly improve trust and performance.

According to Barnes et al. (2011) in military scenarios, uncertainty information should be presented in terms of success rather than failure. This finding guided the study design in presenting information as robot confidence rather than uncertainty. Both uncertainty and confidence convey the same information, just inverted. One limitation identified in the literature is that the combination of explanations and confidence level information were not explored, hence the current proposed study.

PROJECT OVERVIEW

The Present Research

This study looks at how confidence level information impacts trust in human-robot teams by examining the interactions of a human-robot team in a simulated urban search and rescue environment. The goal of this study is to provide a better understanding of how confidence level information can impact team trust in complex task environments, such as an urban search and rescue task, and how that information influences operator decisions and reliance on the robot. Therefore, the following was hypothesized:

Previous research has shown that confidence level information can improve trust in human-robot teams, especially in uncertain and unstable environments (Chen et al., 2014; Chen et al., 2016).

H1: Teams in conditions receiving confidence level information will have a higher level of trust in the robot than conditions in which confidence information is not displayed. The increase of transparency through confidence levels will result in more appropriate reliance.

H2: Teams that receive confidence information will have more appropriate reliance on the robot than teams who do not receive confidence information.

This study is part of a larger effort to study human-robot interactions within a synthetic Urban Search and Rescue (USAR) environment. The synthetic task environment was designed to mimic a building after a disaster. The task was to navigate through the environment, locate victims, identify their medical status, and then choose what necessary resources (first-aid kit, EMT, medical evacuation) to provide to them based on the robot's observations and recommendations. Some of these resources are limited, therefore participants must also manage their resources responsibly. The team was also responsible for overcoming any roadblocks or disruptions that occurred throughout the task. There were three main task roles: Navigator, who was responsible for routing the robot through the environment and making decisions on victim resources; the Robot, who navigated through the environment, scanned victims for biometric data, and recommended which resource to provide; and the Incident Commander, who oversaw assigning resources and communicating feedback on victim status and the environment. The role of Navigator was filled by a participant, whereas the roles of Robot and Incident Commander were filled by trained experimenters acting as confederates.

Previous research on the impact of confidence level information has primarily been in more simple task environments that generally only required binary decisions. Because of the complex nature of this task and team interactions, we can see how confidence levels impact decision making when facing non-binary decisions, and therefore the trust in the robot.

METHODS

Data for this thesis came from an experiment which took place virtually through Arizona State University in 2022.

Participants

Participants for this study consisted of 66 volunteers from the university population. The 3-agent teams were composed of one human participant and two confederates. Participants were randomly assigned to one of three conditions. Participants were compensated $10 per hour.

Apparatus

Testbed and Task. An urban search and rescue synthetic task environment created in Roblox was used for this experiment. This testbed is a team-based task environment in which participants must work with a robot and incident commander to navigate through a building following a disaster and locate victims. This testbed was designed to mimic the cognitive, coordination, and communication demands of a real urban search and rescue team with non-expert participants.

The main goal of this task is to navigate the robot through a destroyed building in order to locate victims and assign them different medical resources. The participant must communicate with the robot to help it navigate through the environment by directing it where to go using a map. The participant must also decide what medical resource each victim needs, based on observations of the environment and recommendations from the robot. The participant can choose to provide the following resources to victims in the environment: no action, first aid kit, emergency medical team (EMT), or medical evacuation (med-evac). Medical resources are limited so they must be correctly assigned to victims. The decision will be communicated to the incident commander via voice chat. The participant will then receive feedback from the incident commander on whether the victim recovered with the resource they assigned to them to provide feedback on correct decisions. The participant may also receive feedback and external information from the incident commander throughout the missions.

The participants are informed during training that the robot is, on average, 70% accurate in its recommendations. In order to hold this constant, the robot makes the correct resource recommendation for victims 70% of the time in each mission. The other 30%, the robot's recommendation is incorrect due to an environmental factor the participant must identify. In the mission with confidence information, the confidence levels average to 70% for the overall mission, matching the accuracy. Individual confidence levels for each victim are randomized.

Training and Conditions. Training for this experiment consisted of an interactive PowerPoint training and a hands-on training mission. During the PowerPoint training, participants learned about their assigned role as the navigator and tasks associated with that role. They were instructed on how to communicate with both the robot and incident commander, how to read victim information, and how to assign victim resources. The hands-on training mission consisted of a short 5-minute mission in the task environment. During this time, participants practiced navigating through the environment with the robot, identifying victims, and assigning resources. They also practiced communicating with their teammates. Experimenters followed a script to walk participants through the mission, and answered any questions regarding tasks or roles to establish understanding. Following the training mission, participants complete a pre-test survey and begin Mission 1.

Design. This experiment is designed as a 3x2 mixed subjects design. There were three between-subject variables in this experiment manipulating the modality of communications between the robot and participant. The three conditions will be *graphics only, text only, and both text and graphics.* Within each condition, there will be a within-subject variable of *confidence information*. Participants will be randomly assigned to receive one of the three modalities for both missions, and will receive *no confidence* *information* for the first mission and will receive *confidence information* in the second mission. In all conditions, participants will receive information about each victim and a recommendation on which resource to assign from the robot. In the first mission, participants will only receive observations and recommendations from the robot through the assigned modality. In the second mission, participants will also receive confidence information from the robot regarding its recommendations and observations through the modality assigned. Participants will be told they are interacting with a different model of the robot in mission two, so that they are not biased to trust it when receiving or not receiving confidence information.

Procedure

Participants were randomly assigned to one of the three conditions and assigned the role of navigator in the team. They were told that the robot is an artificial agent, however a trained experimenter controlled the robot. The role of incident commander was also filled by a trained confederate. Participants participated in this experiment virtually through Zoom, in which they were able to view the robot's camera and screen, use Roblox text chat, and have audio communication. Participants signed an informed consent form and completed a color blindness test to ensure they met the requirements to participate in the study. After that, participants began the training PowerPoint and completed the training mission. Following the training mission, participants completed a pretest questionnaire, including the Propensity to Trust questionnaire. Next participants completed Mission 1, which consisted of 20 possible victims to locate. After Mission 1 was complete, participants completed trust and workload surveys before beginning Mission 2, which also consists of 20 possible victims and has 5 perturbations. Upon completing Mission 2, participants completed another trust and workload survey, along with demographic and engagement questionnaires. Participants were debriefed at the end of the experiment session and were compensated $15 for their participation. Experiments took approximately one and a half hours to complete.

Measures

Performance. Performance will be measured to assess the team as a whole. Performance will be based on a point system and will be calculated by the number of victims located, whether or not the right resource was allocated, and being able to overcome perturbations. The resource needed for each victim is determined by the status of health indicators analyzed by the robot and any external factors. There are times during which the robot recommends a resource for a victim based on its health status, but environmental factors like dust or fire indicate that a higher resource is needed. A participant is able to assign a higher resource than needed to a victim for it to survive, but will be penalized later on when they do not have enough of a given resource for victims who actually need it. Because there is a large number of possible victims, participants must work quickly and effectively with the robot to locate all of them. Each mission is timed, and any victims not yet identified by the participant will result in a loss of points. On the other hand, if a participant finishes early, they will be awarded a bonus point for each 30 second interval left on the timer. Participants will receive a score out of a total number of possible points for each mission.

Subjective Trust Scale. To assess trust, Chancey et al.'s 2017 Trust Questionnaire will be used to collect a subjective assessment of trust following each mission. This questionnaire has three main dimensions: performance, process, and purpose. Each dimension contains five questions related to trust that are answered on a 7-point likert scale. The results from this questionnaire will be scored and compared between missions and conditions.

Participants will also complete a Propensity to Trust questionnaire prior to starting mission 1 in order to establish a baseline of participant trust between other humans and with automation.

Reliance. Reliance will be measured by calculating the total number of decisions made that were the same as the robot recommendations in proportion to the total number of robot recommendations, regardless of whether or not that decision was correct. Reliance will be calculated for each mission across participants.

Other Measures. This experiment also will collect other measures that will not be examined in this study, including situation awareness, workload (NASA-TLX), perceived task interdependence, engagement, and demographics.

ANALYSIS AND RESULTS

An a priori power analysis was conducted using G*Power3 (Faul et al., 2007) to test the difference between the means of 3-conditions by 2-confidence level using an *F-test*, with a large effect size ($\eta_p^2 = 0.08$; Cohen, 1988), and an alpha (α) of 0.05. According to the result, 66 participants are needed to run the experiment.

Trust

A 3x2x2 mixed ANOVA with Information Modality as the between subjects variable and Mission and Confidence as the within-subjects variables was conducted to examine the effects of confidence information on trust scores. The descriptive statistics show that for mission 1 and mission 2, those who received confidence had higher average trust scores in that mission than those who did not (Figure 1). For mission 1, participants who received confidence (M = 5.34) had higher trust than those who did not receive confidence (M = 5.20). The same was true for mission 2, where participants who received confidence (M = 5.14) had higher trust than those who did not receive confidence (M = 5.09). While the descriptive statistics results indicated that there is a slight difference between mean trust score by confidence, the results from this study did not reach significance $F(1, 46) = 2.226, p = 0.121$.

Figure 1.
Average trust score by Mission and Confidence

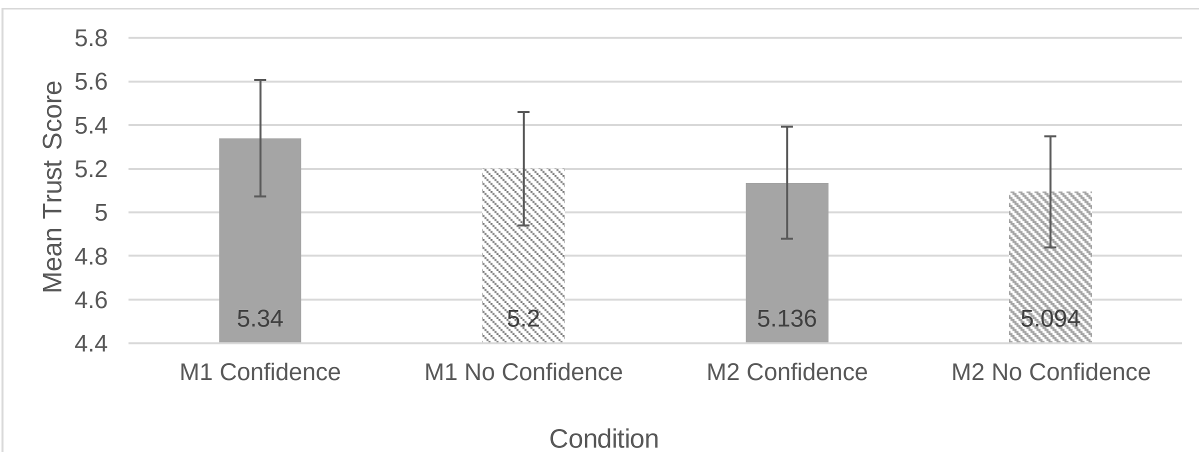

Note. Error bars indicate 95% CI.

Additional unplanned analyses revealed significant interaction between information modality and confidence. Pairwise comparisons revealed that participants in the graphic information modality condition had significantly higher trust than those in the text condition during missions in which they received confidence information (p = .025), but not significantly higher than participants who received both combined text and graphic (full) modalities (p = .027). There was no significant difference in average trust scores between full and text conditions with confidence. This finding indicates that confidence improved trust when the information was presented in a graphic modality, rather than textually (Figure 2).

Figure 2.
Average trust score by Information Modality and Confidence

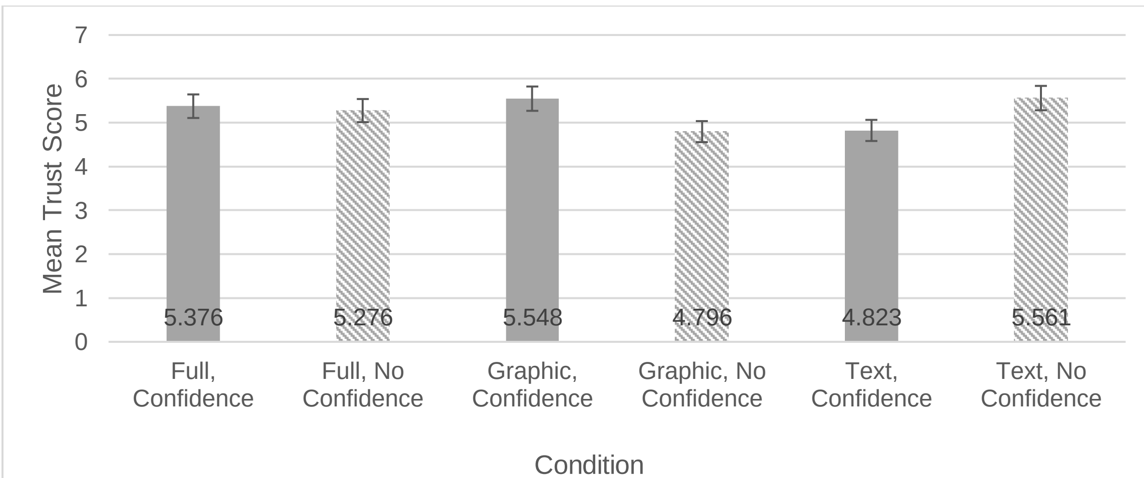

Note. Error bars indicate 95% CI.

To test the hypothesis related to reliance on the robot, reliance for each mission was calculated and an analysis was conducted. Reliance scores were analyzed using a 2 (confidence) x 2 (mission) mixed ANOVA. There were no significant interaction effects revealed. However, proper reliance calibration for this study was said to be at .70. The means for both missions with and without confidence were all extremely close to .70 (Table 1). These results suggest that reliance was properly calibrated for all conditions, regardless of whether confidence information was presented.

Table 1.
Mean Reliance Score by Mission and Confidence

Mission 1 with Confidence	.724
Mission 1 with No Confidence	.701
Mission 2 with Confidence	.715
Mission 2 with No Confidence	.697

DISCUSSION

Summary

This study was motivated by two research questions. First, how does the presence of confidence information impact trust in a human-robot team. Second, how does the presence of confidence information impact reliance calibration on the robot in a human-robot team. These research questions were explored by examining trust in human-robot teams while they received information in one of three modalities: Text, Graphic, or Both, and confidence or no confidence information. These human-robot teams worked together in a virtual synthetic task environment created in Roblox to complete two missions, one in which they received confidence information from the robot and the other in which they did not. After each mission, participants completed a trust questionnaire by Chancey and colleagues (2017) to gather data regarding their subjective trust toward the robot. Reliance was also calculated based on the number of suggestions from the robot that were accepted by the participant, regardless of if the suggestion was correct or not.

Results of this study indicated some differences in mean trust score between missions in which participants received confidence information and missions in which they did not. However, no significant results support the hypothesis that confidence information can increase trust in human-robot teams. This study also showed that regardless of whether confidence information was present or not, reliance was mostly appropriately calibrated. Overall, this study provides insight into how small things such as confidence information and the modality of presentation can influence trust in human-robot teams. However, further research is needed on the impact of such factors on other relationships, such as workload and situation awareness. Further research is also needed to determine if confidence information can be used to appropriately calibrate reliance in human-robot teams.

Discussion of Findings

It was hypothesized that participants in conditions receiving confidence level information will have a higher level of trust in the robot than conditions in which confidence information is not received. However, results for this hypothesis did not reach statistical significance. Another indication that supports the above hypothesis is that participants who received confidence information in Mission 2 had a smaller decrease in trust from Mission 1. Mission 2 of the study was inadvertently harder than Mission 1, which could be why overall mean trust for Mission 2 was lower across both conditions. The smaller decrease in mean trust score by those who received confidence in Mission 2 suggests that the confidence helped with trust in the robot. Additional analyses revealed that information modality combined with confidence information had a significant effect on trust, as participants who received confidence information graphically had significant higher trust than those who received it through a textual modality.

It was also hypothesized that confidence information could be used to appropriately calibrate reliance. Although the results from this study did not show significant effects of confidence on reliance, the means did suggest that reliance was already appropriately calibrated for this study.

Applications

The use of confidence information continues to show potential for improving trust in human-robot teams, especially when using the robot as a decision aid. In future human-robot team efforts, confidence information could be provided when the robot is not 100% reliable. Confidence information helps convey uncertainty to the human operator while also helping to maintain trust in the robot. This communication of uncertainty information helps the user determine if they should rely on the robot's recommendation or not.

Limitations

There were a number of limitations to this study. First, the sample was underpowered. This could have been a factor in why none of the results reached statistical significance. Another major limitation to this study is the difference between missions. The missions were intentionally designed in this experiment to be different to prevent learning effects. However, in attempt to make Mission 2 different than Mission 1, Mission 2 was inadvertently harder than the first mission. There was an unintended mission effect on trust between Mission 1 and Mission 2, where trust was significantly higher overall for Mission 1 than Mission 2. This study was also part of a larger effort to explore the impact of different information modalities on human-robot team interactions, which additional analyses revealed did have some impact on trust. The significant differences in trust by information modality condition also have been impacted by other factors such as workload, as the presentation of information through text could have resulted in higher participant workload compared to the graphic condition. The different presentation types of confidence may have also influenced how participants perceived the robot's confidence, as one was framed textually and the other was framed graphically.

Other limitations are related to the study design of this experiment. This study was conducted in a virtual environment with remote participants, leaving room for technological errors such as latency and screen differences. Participants also answered a multitude of survey questions throughout the experiment, and many times tried to complete the surveys as quickly as possible. This could have resulted in many false reports of perceived trust on the robot. Many participants likely clicked through surveys without completely reading questions, which may have impacted results. In the future, solutions to keep participants engaged should be considered, as well as ways to control for distractions. For example, requiring cameras to be on, or conducting the experiments inside an in-person lab environment.

Lastly, participants were informed prior to the study that the robot was reliable 70% percent of the time in its recommendations. This knowledge of information by the participants may have had a larger impact on trust calibration than confidence information, which is why reliance was appropriately calibrated across all participants and missions. Another explanation is that the task may have been too easy, which meant the participants never really had to rely on the robot for a recommendation. Future studies should look at how increasing workload impacts reliance on the robot. Other future studies should examine if providing information on a robot's reliability helps to truly calibrate appropriate reliance.

Conclusion

Trust has long been the key to understanding human-robot team interactions and ensuring an effective human-robot team. Research has been conducted to examine the impact of both confidence information and information modality on trust in human-robot teams separately, but not together. This study attempts to combine these two variables in order to determine the full impact they potentially have on human-robot teaming. Results from this study can be applied to future uses for human-robot teams where trust is necessary between teammates and time is limited. For example, highly complex and stressful environments such as search and rescue missions. This thesis applied concepts of the Situation Awareness-based Transparency (SAT) model (Chen et al., 2016) and other trust concepts in attempt to better understand the effects of confidence information on human-robot teams. The results from this thesis suggest that confidence information may be used to increase trust in human-robot teams, especially in complex task environments such as an urban search and rescue mission. Differences in trust scores were discovered between missions with and without confidence, although none of which reached significance. However, it was discovered that information modality had a significant effect on trust. This thesis may help influence hypotheses and design considerations for future human-robot team research in a similar context. Future directions surrounding trust and confidence information should make sure to account for any additional factors, such as mission effects or workload, as well as some of the other limitations identified from this study. Future research should continue to focus on understanding how confidence information can impact trust and other measures in human-robot teams, as well as the most effective ways to calibrate appropriate reliance.

REFERENCES

Barnes, M. J.; McDermott, P. L.; Hutchins, S.; Rothrock, L. Framing, Loss Aversion, and Visualization of Risk for a Dynamic Simulation Environment. Journal of Cognitive Engineering and Decision Making 2011, 5 (3), 294–308.

Bass, E. J.; Baumgart, L. A.; Shepley, K. K. The Effect of Information Analysis Automation Display Content on Human Judgment Performance in Noisy Environments. Journal of Cognitive Engineering and Decision Making 2013, 7 (1), 49–65.

Billings, D. R., Schaefer, K. E., Chen, J. Y., & Hancock, P. A. (2012, March). Human-robot interaction: developing trust in robots. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction* (pp. 109-110).

Brown, D. (2021, June 30). Throwable military robots sent to assist with Florida condo collapse.
Retrieved from https://www.washingtonpost.com/technology/2021/06/30/throwable-robot-florida-condo-collapse/

Chancey, E. T., Bliss, J. P., Yamani, Y., & Handley, H. A. (2017). Trust and the compliance–reliance paradigm: The effects of risk, error bias, and reliability on trust and dependence. *Human factors*, *59*(3), 333-345.

Chen, J. Y., Procci, K., Boyce, M., Wright, J., Garcia, A., & Barnes, M. (2014). *Situation awareness-based agent transparency*. Army research lab aberdeen proving ground md human research and engineering directorate.

Chen, J. Y., Barnes, M. J., Selkowitz, A. R., Stowers, K., Lakhmani, S. G., & Kasdaglis, N. (2016, March). Human-autonomy teaming and agent transparency. In *Companion Publication of the 21st International Conference on Intelligent User Interfaces* (pp. 28-31).

Chen, J. Y., Lakhmani, S. G., Stowers, K., Selkowitz, A. R., Wright, J. L., & Barnes, M. (2018). Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical issues in ergonomics science*, *19*(3), 259-282.

Cohen, J. (1998). *Statistical power analysis for the behavioral sciences*. Hillsdale, N.J.: L. Erlbaum Associates.

Faul, F., Erdfelder, E., Lang, A., and Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. doi: 10.3758/BF03193146.

Johnson, C. J., Demir, M., McNeese, N. J., Gorman, J. C., Wolff, A. T., & Cooke, N. J. (2021). The Impact of Training on Human–Autonomy Team Communications and Trust Calibration. *Human Factors*, 00187208211047323.

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human factors*, *46*(1), 50-80

Wang, N., Pynadath, D. V., Rovira, E., Barnes, M. J., & Hill, S. G. (2018, April). Is it my looks? or something i said? the impact of explanations, embodiment, and expectations on trust and performance in human-robot teams. In *International Conference on Persuasive Technology* (pp. 56-69). Springer, Cham.

Wang, N., Pynadath, D. V., & Hill, S. G. (2016, May). The Impact of POMDP-Generated Explanations on Trust and Performance in Human-Robot Teams. In *AAMAS* (pp. 997-1005).