A Mixed Analysis of Influencing Factors for Trust in a Risk-Aware Autonomy

Lixiao Huang, M.L. Cummings, and Masahiro Ono
Arizona State University, Duke University, and California Institute of Technology

To study the factor of controllability for humans' trust in autonomy in decision making, 73 participants used a risk-aware, autonomous planner to navigate an underwater robot in a software-simulated, resource-limited, and risky environment. The experiment examined the association of controllability of the leg size of the path planner with participants' trust, as well as the underlying reasons. The quantitative analysis showed no significant effect of controllability on trust. However, a verbal data analysis method, which systematically coded and quantified participants' reasons for choosing their trust levels at different intervals during the experiment, showed that the dimension execution and control was the most frequently mentioned among four emerging practical dimensions of influencing factors for their trust. The other three included risk evaluation; training and learning; and general attitudes. The findings suggested further research on more aspects of controllability, especially the ones that support personal plans, from a human-autonomy interface design perspective.

INTRODUCTION

While autonomy of different types has been widely researched and applied in significant areas such as aviation, industry, and workplace (Zuboff, 1988), it is suggested to use a moderate level of decision automation in which the human operator is involved in the decision-making process and makes the final decision before the autonomy is 100% reliable and especially when the decisions involve lethality and significant consequences (Parasuraman et al., 2007). To facilitate human–autonomy team performance for high-risk situations with a moderate level of decision automation, important questions arise such as what are the influencing factors for humans' trust in the autonomy from a practical perspective and what could be done to help calibrate their trust?

In particular, risk-aware autonomy makes decisions with explicit considerations of risk and helps humans quantify risks (Ono et al., 2013). A notable approach to risk-aware planning is based on stochastic optimization, where users specify the upper bound of the probability of failure of an automated operation (Majumdar & Pavone, 2020; Ono et al., 2015). However, humans may struggle to understand probabilistic estimates of risk (Tversky & Kahneman, 1973), and their lack of understanding of how autonomy maps risk estimates onto the real world can lead to incorrect mental models of trust that can ultimately hamper system performance. Meanwhile, autonomy may fail to operate reliably in complex, uncertain, and high-risk situations, in which risks are not included in the model, or the probability of risk is difficult or impossible to accurately quantify (Stimpson et al., 2017).

Therefore, appropriate trust in autonomy has a significant impact on human–autonomy team performance in terms of reliance on autonomy in operators' action plans. Trust is defined as a belief that an agent will do what is expected in order to help achieve specific goals (Lee & See, 2004). Overly trusting autonomy may prevent people from intervening when autonomy malfunctions. For one example, the crew in the Royal Majesty shipwreck solely relied on an automated warning, even though the navigation data was erroneous for 24 hours (National Transportation Safety Board, 1997). Over-reliance on Tesla's AutoPilot is another example that has been shown to be deadly (Golson, 2017). Conversely, distrusting autonomy can increase accidents due to inevitable human errors (Zuboff, 1988). For example, fault protection software designed to prevent astronauts from calling wrong programs (e.g., "P01") was rejected by NASA astronauts in 1960s, but later led to the loss of critical data (Mindell, 2008).

To examine trust from a human–autonomy interface design perspective, we look at trust at three different layers, including dispositional trust, situational trust, and learned trust (Hoff & Bashir, 2015). Dispositional trust is a human's tendency to trust or distrust autonomy and is independent of context. Situational trust is a type of trust that depends on the evaluation of the immediate situation and environment. Learned trust is based on humans' knowledge and experience of the past performance of the autonomous system (Lee & See, 2004). These three layers can evolve and change dynamically (Hoff & Bashir, 2015).

Correspondingly, each of the three layers of trust has influencing factors, based on empirical studies (Hoff & Bashir, 2015). Dispositional trust is influenced by culture, personality, gender, age, etc., and thus is less likely to be changed. Situational trust depends on the real-world situation, which is unpredictable sometimes, and thus is difficult to be controlled. Learned trust is based on the system evaluation, which may be specified as learning three aspects of the autonomous system: performance, process, and purpose (Lee & See, 2004). Performance refers to the reliability, predictability, and ability of the autonomy. Process refers to the mechanism of the autonomy (how it works). Purpose refers to whether the intent of the autonomy matches the function. Therefore, given a system that is already built, there is greater chance to adjust people's learned trust than other two layers through the design features of the human–autonomy interface, such as transparency and feedback (Bagheri & Jamieson, 2004) and level of control (Verberne et al., 2012). The design features of the interface may enhance the communication between the autonomy and humans.

In Norman's recurring Seven Stages of Action (Norman, 1988), evaluation and execution are two important processes in humans' interactions with any machine interface. People evaluate the information of a system to form the intention to act. And then they change the state of the task environment through executing their plan with a supporting control panel. Given a known risky situation, controllability is expected to be a critical aspect of interface design that influences human perception and performance (Leotti et al., 2010). Controllability allows humans to be part of the decision making and make the final decision, thus likely influencing humans' calibration of trust.

To understand human–autonomy teaming in high-risk decisions, this study examined influencing factors, especially the factor of controllability, for humans' trust in the risk-aware autonomy that calculates risk probability and suggests short paths in a risky environment. Both the trust ratings and the underlying reasons are important and text analysis has been advocated in trust research (Lee & Kolodge, 2019), so this study explored a mix of quantitative and qualitative analysis.

METHOD

Participants

73 participants were from a southeastern US university (*male N* = 38; *age mean* = 22.8, *SD* = 7.1; *undergrad* = 64.4%, *graduate* = 26.1%, and other = 9.5%; *White* = 41%, *Asian* = 38%, *Hispanic* = 11%, and *Other* = 10%).

Interface of the Software Platform and Experiment Task

The software platform developed for this study is called the Human–Autonomy Interface for Exploration of Risks (HAIER). It was played on a 21-inch screen computer with a mouse. The primary goal of the experiment task was to safely navigate an underwater robot across a risky environment filled with deadly obstacles. To win the bonus incentive of $100, participants needed to achieve the highest performance score possible under four constraints: (a) the shortest path possible should be planned; (b) participants should avoid collisions with the red obstacles, as any collision represented a fatal accident and would end the game with a zero performance score; (c) participants were given a risk budget of 10% for each test session—each leg bears some percentage of risk and the upper bound allowed risk for the whole mission is 10% (Ono, 2012); and (d) the underwater robot needed to surface periodically for positional updates up to 6 times. The performance score is subtracting the portions of overused risk budget, overused surfacing times, and extra path length from a total of 100 points. [A demo of the task is on Youtube] (search "HAIER041717").

In a testing trial, participants first evaluated "how risky the situation is" (not risky, kind of risky, and very risky) on the top right of the interface. Then they used the control panel to choose the risk level they wanted for the path leg. Participants in the No Control group used an interface that had the choice of risk level for that leg but no option of changing an individual leg size. An individual leg size was chosen by the autonomous planner based on participants' selected risk level on the control panel (Figure 1). Participants in the With Control group used an interface with an additional option to select a small, medium, or large leg size. They used the control panel to compare path options on the map and confirmed a choice. Then they answered two pop-up questions regarding their confidence about their last path leg and their trust in the autonomy ("How much do you trust that the autonomous planner generates the best solution?"). The definition of trust was explained as a personal belief that the autonomous planner would do as it is expected to generate the shortest path given your choice of risk level for the leg you are planning.

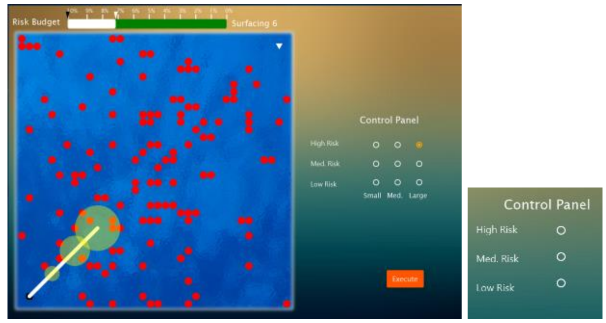

Figure 1. Interface with control of leg size (left), without (right)

A risk budget meant that no more than 10% total in risk of collision should occur for trip on each map, and the risk for every leg planned (the path between two surfacing points) was based on the combined distance to the obstacles. Thus, a participant could choose to spread the risk out over each path leg planned or could use most of the risk to negotiate a single leg. The risk budget was a soft constraint, so overuse would not end the game but reduce overall performance. Similarly, overuse of the surfacing budget would reduce the performance score but not end the game. The yellow circles along the path showed the possible surfacing positions of the robot given a proposed path plan, with growing uncertainty in position predictions as participants planned longer legs.

The autonomous path planner used a *Dijkstra* algorithm to generate the shortest path between the current position of the robot and the destination (i.e., a global planning approach) (Ono et al., 2015). But the execution of the path was broken down into several surfacing points for a position update, which required participants to do local planning for each path leg. HAIER allowed participants to choose among three risk levels (low, medium, and high) per leg. Given a chosen risk level, the autonomous planner displayed a corresponding path option on the map: high-risk paths were shortest possible lines through crowded obstacles and low-risk options tended to take longer paths to stay away from obstacles. Thus, the journey from start to finish included a series of legs between surfacing points. The task was complex because the shortest paths were sometimes riskier because of more obstacles in the center of the map; the detours were less risky, but then the overall path lengths were longer. In addition, due to unpredictable ocean currents, the robot could drift from the planned path when traveling under the water, resulting in a potential collision with nearby obstacles. The shorter the leg the robot took under the water each time, the less uncertainty of its location when surfacing, yet taking too many short legs could prevent the vehicle from finishing within six legs. The whole game was a balance of gains and potential losses in the presence of uncertainty. The unpredictable drifting, together with the requirement of controlling risk budget consumption, surfacing times, and path length, made this a judgment an uncertainty decision process (Tversky & Kahneman, 1973).

Procedure

Recruited from a campus advertisement, the participants were paid $15 for the 60-minute experiment, with a $100 bonus for the top performer. After signing a consent form, a demographic survey, and a risk propensity survey, participants received training (5-10 minutes) prior to each testing trial by hearing the experimenter introduce each element on the autonomy interface and the performance goals, and then practiced path planning until they finish the training map and their questions are answered. Two testing trials started when participants verbally confirmed they felt ready for testing, each for 5-15 minutes. Screen recordings of the last testing trial were replayed for a recall interview, followed by a post-experiment online survey and a debriefing.

Measures of Trust

Trust was measured in two ways. First, after planning each leg but before the resulting surfacing location was revealed, a pop-up question asked participants to report "how much do you trust that the autonomous planner generates the best solution?" on a 5-point Likert scale (1 = no trust at all, 5 = complete trust). Each map took 5 to 8 legs to finish, so participants rated their trust levels 5 to 8 times.

Second, when replaying a screen capture video to provide a specific context for retrospective verbal protocol analysis, the experimenter orally asked participants why they chose the trust level at each leg during the testing session. Following the interview, participants listed their trust strategies in an open-ended essay question in an online survey, "How did you choose your trust [levels] toward the automation?" This question aimed to capture the kinds of influencing factors for trust that participants developed based on their interactions with the autonomy.

The Likert scale data used a t-test for data analysis, and the qualitative written responses were systematically coded, with details in the following section.

Qualitative Data Coding

The analysis of participants' reasons for trust used Geisler's verbal data analysis method (Geisler, 2003). A typical procedure of this method includes (1) segmenting participants' written reasons into T-units (minimal terminable meaning units) (Hunt, 1965), (2) developing mutually exclusive categories of themes for the T-units, (3) having two independent coders code the T-units by judging which category each T-unit falls under, (4) calculating the inter-rater reliability, (5) refining the categories by merging overlapping ones, splitting double meaning ones, or specifying definitions, (6) repeating steps 3-4-5 until the inter-rater reliability is over 90%, and (7) illustrating the categories, their frequency, and sometimes the change of frequency over time. This method transforms text data into quantities and visualizations, helping researchers explore the text content themes, relationships, and change patterns over time. Geisler's verbal data analysis method is consistent with and capitalizes on protocol analysis, content analysis, conversation analysis (Koschmann, 2013; Zemel & Koschmann, 2013). In this study, after five rounds of repeating steps 3-4-5, the final coding had a simple agreement of 97% inter-rater reliability at the dimension level and 93% at the category level, with a *Kappa* = .96.

The segmentation of 73 participants' responses resulted in 155 T-units (No control group = 90, With control group = 65), which then were used to develop the data-driven bottom-up codebook (the hyperlink shows the full codebook) using Saldana's coding method (Saldaña, 2012). Based on the relations of the categories and literature, 16 categories were grouped into five dimensions, as described below.

General attitude (GA). This dimension consists of statements that only describe participants' general trust tendencies that are independent of the context, similar to dispositional trust (Hoff & Bashir, 2015).

Training and learning (TL). This dimension consists of responses about becoming familiar with the autonomy: understanding the algorithms and definitions, learning the past performance of the system, and gaining experiences with the autonomy. Specifically, participants mentioned the length of interaction experiences, mechanisms of how the autonomy works, and overall impressions of the autonomy's past performance. These factors pertain to adequate knowledge of the autonomy and can be enhanced accordingly to potentially induce informed trust. Therefore, they are grouped under the dimension of training and learning.

Execution and control (EC). This dimension includes categories pertaining to participants' action plans with the autonomy, and how human–autonomy interface supports participants in carrying out their action plans. The choice of risk levels and leg sizes are two major control options in HAIER, and their combinations result in different path options on the map. Therefore, categories that are related to the choice of risk levels, leg sizes, and path options, are grouped under the dimension of execution and control. The choice of risk levels is associated with the appearance of proposed paths, their comparisons with users' envisioned routes, and the remaining risk budget. The choice of path options sometimes involves evaluation statements about the paths, including specific aspects (e.g., the direction of the path and leg size) and overall or emotional comments regarding the paths (e.g., "the generated path was strange or frustrating").

Risk evaluation (RE). All statements related to risks are grouped under the dimension of risk evaluation due to the diverse statements about risks in the data and the salient role of risk in trust in autonomy in the literature (Hoff & Bashir, 2015; Stimpson et al., 2017). The dimension of risk evaluation is related to situational trust because of its focus on the environment and context (Hoff & Bashir, 2015), specifically, deviations from planned paths, size of yellow circles regarding uncertainty in future surfacing locations, number of obstacles, and overall perception of danger. The dimensions of EC and RE are closely related because risk evaluation accompanied path evaluation and influences path choice. The categories mentioned under this dimension mainly illustrate what participants perceive as risk and what influence their risk perception (Norman, 1988).

Other. Only one response went astray from the question of whether this participant trusted the autonomous planner or what reasons influenced the trust and was coded as other.

RESULTS

Quantitative Data Analysis

Since each map took at least five legs to finish and N started to decrease at leg six, GLM repeated ANOVA used mean trust ratings per leg for the first five legs as a within-subject variable and Control as a between-subject variable, and the results was statistically insignificant ($F(1,132) = .05$, $p = .83$). Then we averaged all 5-8 ratings of trust rating to get one mean score for each map (i.e., the sum of the 5-8 trust ratings divided by the number of legs for that map) to do a t-test on Control, the difference in mean trust between No Control ($M = 3.37$, $N = 74$, $SD = .79$) and With Control ($M = 3.70$, $N = 60$, $SD = .78$) was also insignificant ($t(132) = .16$, $p = .88$).

Qualitative Results: Verbal Data Analysis

Using a bottom-up approach, four dimensions emerged from the verbal responses of people's strategies in choosing their trust in autonomy and are listed in order of decreasing frequency: execution and control (EC), risk evaluation (RE), training and learning (TL), and general attitudes (GA). Figure 2 shows the frequency of each dimension. At a more specific category level, Figures 5-8 show the frequency of categories in each dimension.

For the dimension of GA (*No Control* = 13, *With Control* = 5, *total* = 18), Figure 3 shows that people in both groups indicated trust and distrust in autonomy. But none in the With Control group said that they were unsure, maybe because the leg size control helped reduce this uncertainty.

For the dimension of TL (*No Control* = 15, *With Control* = 9, *total* = 24), people in both groups showed similar focus on the planner's past performance and their comprehension of planner's mechanism (Figure 4). Only participants in the No Control group reported that interaction length/familiarity with the system is an influencing factor with the system.

For the dimension of RE (*No Control* = 21, *With Control* = 14, *total* = 35), Figure 5 shows that deviation, probability circle, obstacles, and distance to the target are all potential risk elements. People were most concerned about robots' deviation from the planned path, especially in the With Control group. For the dimension of execution and control (*No Control* = 40, *With Control* = 36, *total* =76), Figure 6 shows that the highest number of participants commented about whether the proposed options matched personal ideal plans.

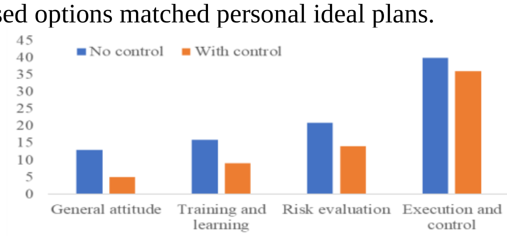

Figure 2. Frequency of dimensions by groups

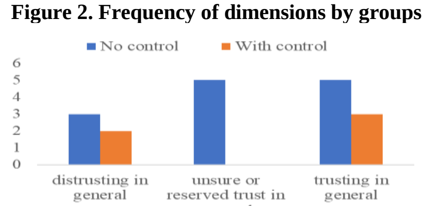

Figure 3. Frequency of categories in the dimension of general attitudes

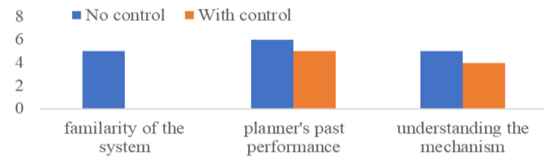

Figure 4. Frequency of categories in the dimension of training/learning

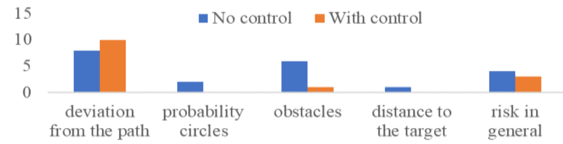

Figure 5. Frequency of categories under the dimension of risk evaluation

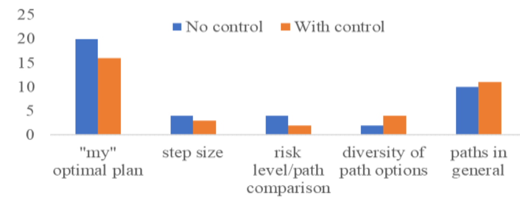

Figure 6. Frequency of categories in the dimension of execution/control

DISCUSSION

This study contributes to trust literature in three innovative ways. First, instead of using narrative imaginary scenarios, a software testbed was developed to provide contextual experience for better validity, in which participants interacted with an autonomous path planner to complete a path planning task with interval performance feedback. Second, instead of using a nondirected and overall trust scale one time, this testbed used a pop-up question to measure context-specific trust multiple times throughout the experiment. Third, this study delved into the reasons for trust ratings by using a bottom-up data-driven approach to systematically code and quantify participants' verbal responses on their trust strategies toward the autonomy, allowing for insight into additional germane influencing factors and participants' interpretation of controllability.

Trust was measured dynamically in two ways: Pop-up Likert-based trust level at each interval and a qualitative written summary of the strategies used for their trust ratings. The Likert data analysis showed that the controllability of the leg size was not a statistically significant factor for trust in the autonomous planner. However, the qualitative data analysis of the written trust strategies produced four dimensions of factors: execution and control, risk evaluation, training and learning, and general attitudes. The dimension of EC was shown to be the most critical in participants' awareness regarding trust, based on its high frequency. Specifically, the category whether system options matched personal plans was the most critical factor in this dimension (Figure 6). A possible explanation for the insignificance of controllability in the quantitative analysis is that people may want more controllability than just controlling the leg size. The follow-up qualitative data analysis indicated that the participants wanted the type of control for the path planner to generate options matching their personal plans and results. This insight is important since trust may be built for these types of interactions by the autonomy aligning with human expectations, not just good performance. This need for more research on the role of user preferences and user interfaces on trust in autonomy (Hoff & Bashir, 2015). Also, it is critical to understand fundamental expectations as well as what can be done to move humans away from biased or incorrect expectations through appropriate training and learning.

The dimension of EC differs from the influencing factor of system performance because EC focuses on whether the human–autonomy interface provides supporting controls for humans to take actions, while system performance focuses on the result alone. Control of actions is an important means to achieve desired human–autonomy teaming performance.

This study also demonstrated that design features for risk evaluation and risk perception may be critical for calibrating trust, and specifically on the visual aids for risk evaluation (Spiegelhalter, 2017), such as the risk budget bar and the yellow circles estimating possible future vehicle locations. High risk perception may relate to low trust in autonomy when controlling a robot on Mars (Stimpson et al., 2017). Research on risk representation has the potential to fill the gulf of evaluation (Norman, 1988) in human–autonomy interface design. Risk representation through quantified risk budget (Ono, 2012) is novel and require additional research on its influence on risk perception and desired controllability.

Training and learning about systems' process, performance, and purpose of a system (Lee & See, 2004) are not the system per se, but a procedure that is separate from these aspects of the system and may also help calibrate trust. Working with a given system may be the case in many human–autonomy interaction scenarios in the real world. When participants came, the HAIER system development was finished. What can be changed is to facilitate participants' learning of the autonomy's performance and mechanism, as well as practicing to become better at working with the autonomy, preferably considering each individual's knowledge and skill base.

Though measuring an individual's trust level toward a system has been studied widely (Hancock et al., 2011; Jian et al., 2000), knowing the underlying reasons for trust and implementing strategies to calibrate trust is equally important, or even more so, but less studied (Lee & Kolodge, 2019). This study found both context-specific and generalizable reasons for trust. Context-specific reasons (e.g., lack diversity of path options) are necessary to improve the specific design of the human-autonomy interface to facilitate human–autonomy teaming (DeCostanza et al., 2018). The generalizable dimensions were practical reasons for the trust that could be implemented through personnel selection, training and learning of the system, risk information display, and controllability of the actions on the interface. Guided by general findings, context-specific and individualized design (DeCostanza et al., 2018) may be explored further for human-agent teaming with moderate level design autonomy.

ACKNOWLEDGEMENTS

We thank software developers Yuansong Feng and Kaijie Chen, the second coder Meghna Mandava and ONR (grant #: N0001417IP00043).

REFERENCES

Bagheri, N., & Jamieson, G. A. (2004). The impact of context-related reliability on automation failure detection and scanning behaviour. *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No.04CH37583)*, *1*, 212–217 vol.1. https://doi.org/10.1109/ICSMC.2004.1398299

DeCostanza, A. H., Marathe, A. R., Bohannon, A., Evans, A. W., Palazzolo, E. T., Metcalfe, J. S., & McDowell, K. (2018). *Enhancing Human–Agent Teaming with Individualized, Adaptive Technologies: A Discussion of Critical Scientific Questions*. US Army Research Laboratory Aberdeen Proving Ground United States.

Geisler, C. A. (2003). *Analyzing streams of language: Twelve steps to the systematic coding of text, talk, and other verbal data*. Longman.

Golson, J. (2017, January 19). *Driver in fatal Tesla Autopilot crash had seven seconds to take action*. The Verge. https://www.theverge.com/2017/1/19/14326604/tesla-autopilot-crash-driver-seven-seconds-inattentive-nhtsa

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y., De Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, *53*(5), 517–527.

Hoff, K., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, *57*(3), 407–434.

Hunt, K. W. (1965). *Grammatical Structures Written at Three Grade Levels. NCTE Research Report No. 3.* http://eric.ed.gov/?id=ED113735

Jian, J.-Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, *4*(1), 53–71.

Lee, J. D., & Kolodge, K. (2019). Exploring trust in self-driving vehicles through text analysis. *Human Factors*, 0018720819872672.

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, *46*(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Leotti, L. A., Iyengar, S. S., & Ochsner, K. N. (2010). Born to choose: The origins and value of the need for control. *Trends in Cognitive Sciences*, *14*(10), 457–463.

Majumdar, A., & Pavone, M. (2020). How should a robot assess risk? Towards an axiomatic theory of risk in robotics. In *Robotics Research* (pp. 75–84). Springer.

Mindell, D. A. (2008). *Digital Apollo: Human and machine in spaceflight*. The MIT Press,.

National Transportation Safety Board. (1997). *Grounding of the Panamanian Passenger Ship ROYAL MAJESTY on Rose and Crown Shoal near Nantucket, Massachusetts* (NTSB/MAR97/01).

Norman, D. A. (1988). *The psychology of everyday things.(The design of everyday things)*. Basic Books.

Ono, M. (2012). *Robust, goal-directed plan execution with bounded risk* [Massachusetts Institute of Technology]. https://dspace.mit.edu/handle/1721.1/71451

Ono, M., Pavone, M., Kuwata, Y., & Balaram, J. (2015). Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, *39*(4), 555–571.

Ono, M., Williams, B. C., & Blackmore, L. (2013). Probabilistic planning for continuous dynamic systems under bounded risk. *Journal of Artificial Intelligence Research*, *46*, 511–577.

Parasuraman, R., Barnes, M., Cosenzo, K., & Mulgund, S. (2007). *Adaptive automation for human-robot teaming in future command and control systems*. Army research lab aberdeen proving ground md human research and engineering ….

Saldaña, J. (2012). *The coding manual for qualitative researchers* (2 edition). SAGE Publications Ltd.

Spiegelhalter, D. (2017). Risk and uncertainty communication. *Annual Review of Statistics and Its Application*, *4*, 31–60.

Stimpson, A. J., Tucker, M. B., Ono, M., Steffy, A., & Cummings, M. L. (2017). Modeling risk perception for mars rover supervisory control: Before and after wheel damage. *2017 IEEE Aerospace Conference*, 1–8. https://doi.org/10.1109/AERO.2017.7943871

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*(2), 207–232.

Verberne, F. M. F., Ham, J., & Midden, C. J. H. (2012). Trust in Smart Systems: Sharing Driving Goals and Giving Information to Increase Trustworthiness and Acceptability of Smart Systems in Cars. *Human Factors*, *54*(5), 799–810. https://doi.org/10.1177/0018720812443825

Zuboff, S. (1988). *In the age of the smart machine: The future of work and power*. Basic books.