Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 20–30
https://doi.org/10.54941/ahfe1008237

Automation Transparency Communication Does Not Improve Supervisory Control Performance, But Shapes Operator Trust

Kaylin Strong[1], Bailee Smith[2], Jill Talley Shelton[3], Shannon McGarry[1], Joseph Coyne[1], Michael Kukenberger[4,5], and Noelle Brown[1]

[1]U.S. Naval Research Laboratory, Washington, DC, USA
[2]Louisiana State University, Baton Rouge, LA, USA
[3]University of Tennessee, Chattanooga, TN, USA
[4]Naval Aerospace Medical Institute, Pensacola, FL, USA
[5]University of New Hampshire, Durham, NH, USA

ABSTRACT

With increased reliance on automated systems in supervisory control environments, such as unmanned aerial vehicle (UAV) missions, it is critical to understand what factors influence the human operator's ability to interact effectively with these systems. The present study was designed to examine whether messages about potential problems with the automation for certain UAV sensor feeds influenced human operator behaviors and trust approach between systems. Two hundred and twenty-one participants, consisting of both Naval aviation enlistees and officers, engaged in a multi-UAV supervisory control task for 28 minutes. They were asked to monitor automated target detection in three UAV sensor feeds, which varied in reliability unbeknownst to the participants, while completing other tasks. All participants received communication about potential problems with two of the sensor feeds after completing seven minutes of the supervisory control task. Subjective assessments of automation trust were administered for each sensor feed immediately before the transparency communication and again at the very end of the task. After the UAV supervisory control task, individual difference measures of cognitive ability were collected. Results revealed that all participants showed some awareness of differences in system performance, but higher cognitive ability was associated with a more distinct pattern of trust over time. These findings demonstrate the importance of considering operator characteristics when designing transparency interventions and supervisory control systems.

Keywords: Human-automation interaction, Trust calibration, Individual differences, Transparency communication

INTRODUCTION

In military supervisory control environments, operators are frequently responsible for monitoring multiple automated systems simultaneously while intervening only when automation fails (Endsley, 2017). Although automation can reduce workload and improve performance, these benefits depend on the operator's ability to maintain effective oversight and detect

Received November 7, 2026; Revised November 24, 2026; Accepted November 27, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

automation errors. Sustained monitoring, however, is limited by declines in vigilance over time, making supervisory control particularly vulnerable to automation failures (Mackworth, 1948).

Successful human–automation interaction depends on automation reliability and how operators develop and calibrate trust in automated systems. Trust influences reliance decisions, monitoring behavior, and the willingness to intervene when automation produces incorrect outputs (Mishler & Chen, 2023). When supervising multiple automated systems, operators often fail to differentiate between systems with varying levels of reliability and instead adopt a *system-wide trust* (SWT) strategy in which all systems are treated similarly regardless of their objective performance (Foroughi et al., 2019; Keller & Rice, 2010; Lopez & Pak, 2020; Rice & Geels, 2010; Walliser et al., 2016). This generalized trust strategy has been associated with inefficient monitoring behaviors and increased workload because operators distribute their attention relatively equally across systems rather than allocating resources according to the system's objective reliability (de Visser et al., 2016).

An alternative approach, *component-specific trust* (CST), proposes that trust should be calibrated independently for each automated system according to its reliability. Under a CST strategy, operators allocate greater attention to less reliable automation while reducing attention on highly reliable systems, thereby improving monitoring efficiency. Although CST is theoretically optimal, evidence suggests that operators rarely adopt this strategy spontaneously, even when system reliabilities differ substantially (Lopez & Pak, 2020; Walliser et al., 2023). Consequently, identifying individual differences and interventions that promote appropriate trust calibration remains an important challenge in supervisory control research.

Automation transparency has emerged as one promising approach for improving human–automation interaction. Transparency refers to communicating information about the automation's processes, limitations, uncertainty, or expected performance to help operators develop accurate mental models of system behavior (Mercado et al., 2016). Providing accuracy information was shown to reduce algorithm aversion (Dietvorst et al., 2018; Yin et al., 2019), a decrease in trust and reliance following algorithm errors (Dietvorst et al., 2015). Meta-analytic evidence indicates that transparency generally improves operator performance, trust, and understanding of automated systems (Sargent et al., 2023). However, transparency is not uniformly beneficial. Devlin et al. (2022) found nuanced benefit to performance with systems of varied reliability attributed to automation transparency. Specifically, automation error detection improved for the system mentioned in the transparency communication but decreased for the system that was not mentioned. Ling et al. (2024) demonstrated that providing visual explanations improved trust and detection performance, whereas additional confidence information sometimes impaired performance. Similarly, conflicting information can increase human workload and perceived stress, and how information is presented can affect trust and comprehension (Balfe et al., 2015; Pareek et al., 2024). Thus, there may be limitations to the benefits of automation transparency.

Research has also emphasized the importance of operator characteristics in supervisory control. Cognitive abilities including working memory, attentional control, spatial ability, and fluid intelligence similarly predict monitoring and multitasking performance (Brady et al., 2022; Redick et al., 2016).

Despite growing evidence supporting the benefits of automation transparency, relatively little research has examined whether transparency can improve monitoring behavior in environments requiring concurrent supervision of multiple automated systems with different reliabilities. Existing studies have largely focused on single automation systems. To extend the current findings, we examined whether automation transparency could improve automation error detection and facilitate the adoption of a component-specific strategy in a multi-UAV supervisory control task. We also investigated the role of cognitive ability in the adoption of trust approaches and in the efficacy of the automation transparency.

BACKGROUND AND STUDY PURPOSE

Prior findings suggested automation transparency benefits may be limited to the system the alert references (e.g., Devlin et al., 2022). Our prior work also suggested the effects on automation error detection may be further limited by a limited range in task performance, and by practice effects (McGarry et al., under review). To address these limitations, we increased the rate of automation errors and presented the uncertainty communication earlier compared to our previous work (Devlin et al., 2022; McGarry et al., 2023; McGarry et al., under review). Despite these changes to the experimental paradigm, there was still little evidence the automation transparency affected automation error detection beyond practice. Thus, the scope of this paper focuses on the relationship between cognitive ability and automation trust. Specifically, we investigate the following: Does cognitive ability moderate the trust approach across varied automation reliability? If it does, is this a consistent effect over time?

METHOD

Two hundred twenty-one U.S Navy and Marine Aviation trainees from the Naval Air Station in Pensacola, Florida participated in the study. Participants were student pilots, flight officers or enlisted trainees assigned to an aviation support rating awaiting training. The majority of participants were male (89.4%) and reported a mean age of 22.8 years (SD = 4.0 years). Groups of 10-20 participants completed the study simultaneously at their own workstation with the facilitation and oversight of a research administrator. The study was approved by the Institutional Review Board at the U.S. Naval Research Laboratory.

Tasks

The Supervisory Control Operations User Testbed (SCOUT™) was created by scientists at the Naval Research Laboratory (Sibley et al., 2016) to

simulate a supervisory control environment for UAV pilots. For this task, participants were asked to monitor three UAV sensor feeds while engaging in other activities. The SCOUT™ interface, as seen in Figure 1, includes a map showing where the vehicle is heading, objects detected by the automation, and ground control intel. Participants were asked to monitor automated target detection activity in the three sensor feeds. Different target objects (triangles, circles, or squares) were designated for each feed, with all object types appearing across feeds. All objects waterfalled down the screen for 14 seconds with targets appearing every 6 seconds and distractors appearing every 5 seconds. All objects had a stimulus onset asynchrony of (+/−) 2 seconds. Participants were asked to monitor the automation's target detection activity in all feeds and to correct automation errors. Automation target detection was indicated by a gold box surrounding an object. The automation could either successfully detect the target (hit), fail to detect targets (miss), detect an erroneous object (false alarm), or correctly reject erroneous objects (correct rejection). Participants were instructed to correct automation misses and false alarms. Notably, the sensor feeds varied in reliability, with the left feed depicting 70% reliability, the middle feed 100% reliability, and the right feed 95% reliability in target detection.

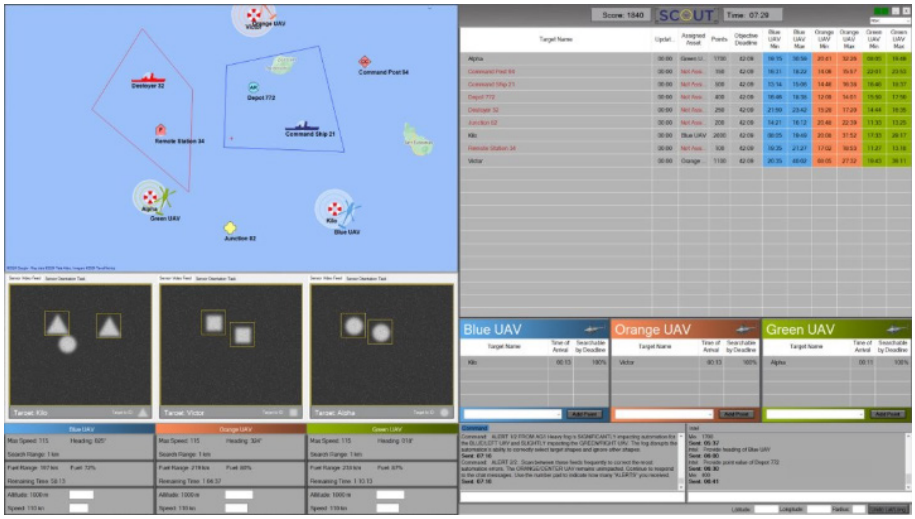
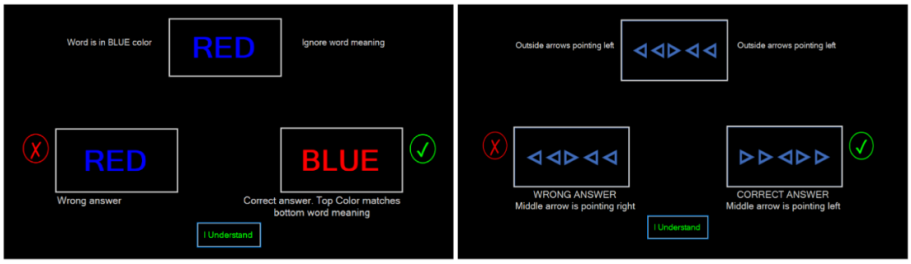


Figure 1: The supervisory control user testbed (SCOUT™).

To assess attention control, participants completed a mixed version of the attention control conflict tasks described in Burgoyne et al. (2023), termed Mixed AC here. In this task, participants saw alternating stroop and flanker trials, but unlike the standard tasks, there is conflict between the cognitive processes required for the stimulus and response options (see Figure 2). Specifically, stroop trials required participants to respond to the font color in the top stimulus, ignoring its word meaning, and to select the matching word meaning, ignoring the font color, among the bottom response options. Red and blue were the only words and font colors presented in the

task. For flanker trials, both the top stimulus and response options showed a row of 5 arrows. Participants were to select the response option whose center arrow matched the direction of the outer arrows presented in the top stimulus. Participants first practiced each task type separately, followed by a brief mixed set of trials. There were two, 90 second experimental blocks of trials presented back-to-back. The entire task took about 7 minutes.



2a 2b

Figure 2: (a) Example of a practice stroop trial in the Mixed AC task. (b) Example of a practice flanker trial in the Mixed AC task.

To assess fluid intelligence, we administered two pattern matching tasks: Ravens Advanced Progressive Matrices (RAPM) and number series task (Raven, Raven & Court, 1998; Redick et al., 2013; Thurstone, 1938). In the RAPM, participants were presented with a set of geometric designs with a related pattern. Participants were to select the geometric design that completed the pattern from 8 response options; answering as many sets as possible within 10 minutes. In the number series task, participants were shown a set of numbers. They were instructed to determine the mathematical pattern that related the number set and select the number that completed the series from 5 response options, answering as many series as possible within 4.5 minutes.

Procedure

Participants were first provided informed consent. Next, participants completed a demographic survey, eye tracker set up and baseline pupillary assessment followed by a PowerPoint presentation of the SCOUT™ instructions as a group. A research administrator addressed any outstanding questions to the group. All participants then completed 7 minutes of the SCOUT™ task, followed by a break in which they completed subjective measures of fatigue (1 (fully alert) to 7 (completely exhausted), max and average workload (1 (nothing to do, no system demands) to 7 (overloaded, system unmanageable) and automation trust. The trust questionnaire consisted of four questions asking participants to provide responses to each question for each sensor feed separately. Specifically, participants were asked to rate their level of trust in the automation aid, the extent of reliance on the automation aid, their self-confidence in performing without the automation aid, and extent that automation improved their performance from 0 (Not at

all) to 100 (Extremely). Following the surveys, the SCOUT™ task resumed and participants received communication regarding potential problems with the reliability of the left and right sensor feeds. The alert messages read as:

ALERT 1/2: Heavy fog SIGNIFICANTLY impacting automation for the BLUE/LEFT UAV and SLIGHTLY impacting the GREEN/RIGHT UAV. The fog disrupts the automation's ability to correctly select target shapes and ignore other shapes.

ALERT 2/2: Scan between these feeds frequently to correct the most automation errors. The ORANGE/CENTER UAV remains unimpacted. Continue to respond to the chat messages. Use the number pad to indicate how many "ALERTS" you received.

Following the alerts, participants completed 21 minutes of the SCOUT™ task and then completed the final iteration of the subjective measures of fatigue, workload and automation trust again. The three cognitive ability assessments were completed thereafter in the following order: Mixed AC, RAPM, and number series.

RESULTS

Prior to analysis, the dataset was screened for completeness, validity, and engagement quality. Participants with incomplete demographic or experimental data were removed from the dataset. Additionally, trust questionnaire response variability screening and missing data criteria were applied to identify careless response patterns. Participants demonstrating little to no variation across subjective ratings ($SD < 3$ points) were flagged for individual case review. Flagged cases were evaluated against secondary metrics, such as performance at or below chance levels across other tasks. In total, 15 participants were removed for a final sample consisting of 206.

Further preparation of the data was required for analysis. To derive a single, comprehensive measure of cognitive ability, an exploratory factor analysis was conducted on the three cognitive task scores. Parallel analysis and scree plots were examined to determine the optimal number of factors to extract. Both methods indicated a single-factor solution. Consequently, a principal axis factor analysis with one factor was performed using pairwise complete observations to handle missing data. Bartlett's method was used to determine factor scores for each participant (Bartlett, 1937). The resulting score represents a latent general cognitive ability factor score and was used as a continuous measure of cognitive ability during exploratory analyses. To assess subsequent group-level comparisons, cognitive ability factor scores were binned into three categorical groups (low, mean, high) based on standard deviation values above, below, and within half a point of the sample mean. Because the standardized factor scores were normally distributed with a mean near zero, participants were categorized into low (24.4%), mean (38.8%), and high (36.8%) groups to establish meaningful separation while preserving balanced distributions and statistical power.

A three-way mixed-design ANOVA was conducted to analyze the effects of cognitive ability on the subjective trust ratings across feeds of varied automation reliability. Time (trust ratings at T1 and T2) and Feed (70%,

95%, or 100% reliable) were within-participants variables with cognitive ability group (low, mean or high) as the between-participants variable. This omnibus test revealed a significant three-way interaction ($F(4,401) = 3.24$, $p = .01$, $\eta^2 = .006$) confirming that subjective trust ratings differed over time depending on both the system reliability and operator cognitive ability illustrated in Figure 3.

To further isolate the source of the three-way interaction, we ran separate two-way repeated measures ANOVA for each level of cognitive ability to evaluate the simple main effects of feed reliability, time, and their interaction. For the low cognitive ability group (orange lines) the Time × Feed interaction was not significant ($p = .12$). However, a significant simple main effect of feed was observed ($F(2, 201) = 8.71$, $p < .001$). To determine where these differences existed, paired-samples t-tests were conducted on average trust ratings collapsed across time for the low cognitive ability group. Results indicated that trust in the 70% reliable feed ($M = 44.53$, $SD = 14.03$) was significantly lower than trust in both the 95% reliable feed ($M = 54.47$, $SD = 11.84$; ($t(48) = –4.0497$, $p < .001$, $d = .579$) and 100% reliable feed ($M = 54.86$, $SD = 12.10$; ($t(48) = –3.8813$, $p < .001$, $d = .554$). Trust ratings did not significantly differ between the 95% and 100% reliable feeds ($t(48) = –0.2729$, $p = 0.786$, $d = .0390$).

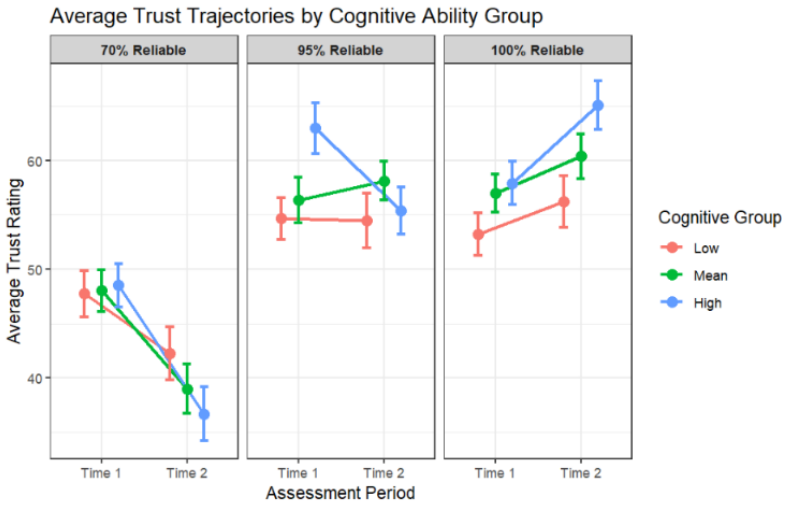


Figure 3: Mean trust scores after the first seven minutes in the task (Time 1), before the transparency communication, and again after an additional twenty-one minutes in the task (Time 2) plotted separately for each cognitive ability group.

The simple Time × Feed interaction was significant for both the mean cognitive ability group (green lines) ($F(2, 201) = 11.65$, $p < .001$) and the high cognitive ability group (blue lines) ($F(2, 201) = 19.62$, $p < .001$). Also, the simple main effect of Feed for both the mean ability group ($F(2, 201) = 34.23$, $p < .001$) and the high ability group ($F(2, 201) = 34.82$, $p < .001$) was significant. Additionally, there was a significant simple main effect of

Time for the high ability group ($F(2, 201) = 7.72$, $p = .006$), but it was not significant for the mean ability group ($F(2, 201) = 1.56$, $p = .208$).

To further analyze the differences between groups, we tested the simple interaction contrast examining the change in trust (T2-T1) for each feed within each level of cognitive ability. The analysis revealed distinct trust approaches among the groups. Trust ratings decreased significantly over time exclusively for the 70% reliable feed for both the low ($t(201) = –2.047$, $p = .042$, $d = .361$) and mean ($t(201) = –4.530$, $p < .001$, $d = .467$) cognitive ability groups. Neither the low nor the mean ability groups demonstrated a significant change in trust over time for the 95% or 100% reliable feeds. Conversely, the high cognitive ability group demonstrated significant changes in trust for all three sensor feeds. Over time, trust for both the 70% reliable feed ($t(201) = –4.936$, $p < .001$, $d = .560$) and the 95% reliable feed ($t(201) = –3.017$, $p = .003$, $d = .370$) significantly *decreased* for this ability group. Additionally, this group exhibited a significant *increase* in trust over time for the 100% reliable feed ($t(201) = 3.263$, $p = .002$, $d = .378$).

DISCUSSION

In this study, we examined the relationship between cognitive ability and subjective reports of automation trust for systems of differing reliabilities. We further explored whether this relationship changed over time. The results revealed that trust in automation did change over time and that this change varied with both the system's reliability and the operator's cognitive ability. Specifically, reported trust was the lowest for the system with the lowest reliability. Notably, trust in this system also declined with experience in the task for all three groups, suggesting a heightened awareness of its poor performance compared to other systems. A limitation of the study is that time on task cannot be divorced from the effect of the automation transparency. That is, the decrease in trust could have been influenced by the alert messages received indicating disruptions to the automation's ability to detect targets, and not simply a mechanism of experience in the task alone.

Though it appeared clear which system the groups did not trust, their trust in the high performing feeds (95% reliable and 100% reliable) differed with cognitive ability. While low and mean ability groups appropriately decreased their trust over time in the least reliable system, they demonstrated no significant change in trust for the other feeds. However, the high ability group demonstrated a significant change in trust across all three feeds. Most interesting was their ability to distinguish the nuanced reliability of the 95% reliable feed, evidenced by their decrease in trust of that system over time.

Taken together, these findings suggest that cognitive ability, specifically fluid intelligence and attention control, may play an important role in how operators calibrate trust across automated systems with differing levels of reliability. While all participants showed some awareness of differences in system performance, higher cognitive ability was associated with a more nuanced pattern of trust over time, with trust decreasing for less reliable systems and increasing for the most reliable system. This pattern is consistent

with a component-specific approach to trust and suggests that individual differences in cognitive ability may influence how effectively operators integrate information about automation performance into their trust judgments. Such findings highlight the importance of considering operator characteristics when designing transparency interventions and supervisory control systems.

ACKNOWLEDGMENT

This research was funded by the Office of Naval Research. The views contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of NRL or the U.S. Government. This work was also supported by the Naval Research Enterprise Internship Program and the ONR Summer Faculty Program.

REFERENCES

Balfe, N., Sharples, S., & Wilson, J. R. (2015). Impact of automation: Measurement of performance, workload and behaviour in a complex control environment. *Applied Ergonomics*, *47*, 52–64. https://doi.org/10.1016/j.apergo.2014.08.002

Bartlett, M. S. (1937). The statistical conception of mental factors. British Journal of Psychology, 28(1), 97–104.

Brady, M. F., Carretta, T. R., Shermadou, Z. S., & Gibson, A. (2022). Aptitude requirements for human operators in human–automation interaction: A meta-analytic review. *Aviation Psychology and Applied Human Factors, 12*(1), 50–62. https://doi.org/10.1027/2192-0923/a000210

Burgoyne, A. P., Tsukahara, J. S., Mashburn, C. A., Pak, R., & Engle, R. W. (2023). Nature and measurement of attention control. *Journal of Experimental Psychology: General*, *152*(8), 2369–2402. https://doi.org/10.1037/xge0001408.supp

de Visser, E. J., Monfort, S. S., McKendrick, R., Smith, M. A. B., McKnight, P. E., Krueger, F., & Parasuraman, R. (2016). Almost human: Anthropomorphism increases trust resilience in cognitive agents. *Journal of Experimental Psychology: Applied, 22*(3), 331–349.

Devlin, S. P., Brown, N. L., Drollinger, S. M., Neilson, B. N., Foroughi, C. K., Sibley, C. M., & Coyne, J. T. (2022). The effectiveness of uncertainty communication with varied automation reliability: Specificity is necessary but (potentially) not enough. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *66*(1), 555–559. https://doi.org/10.1177/1071181322661295

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1):114–126. https://doi.org/10.1037/xge0000033

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3):1155-1170. https://doi.org/10.1287/mnsc.2016.2643

Endsley, M. R. (2017). From here to autonomy: Lessons learned from human–automation research. *Human Factors, 59*(1), 5–27.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics, 16(1),* 143–149.

Fisher, C. R., Morris, M. B., Stevens, C. A., & Swan, G. (2024). The role of individual differences in human-automated vehicle interaction. *International Journal of Human-Computer Studies, 185*. https://doi.org/10.1016/j.ijhcs.2024.103225

Foroughi, C. K., Sibley, C., Brown, N. L., Rovira, E., Pak, R., & Coyne, J. T. (2019). Detecting automation failures in a simulated supervisory control environment. *Ergonomics, 62*(9), 1150–1161.

Keller, D., & Rice, S. (2010). System-wide versus component-specific trust using multiple aids. *The Journal of General Psychology, 137*(1), 114–128.

Ling, S., Zhang, Y., & Du, N. (2024). More Is Not Always Better: Impacts of AI-Generated Confidence and Explanations in Human-Automation Interaction. *Human factors*, *66*(12), 2606–2620. https://doi.org/10.1177/0018720824 1234810

Lopez, J., & Pak, R. (2020). Does one bad machine spoil the bunch? A review of trust in multiple-component systems. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 64, pp. 1546–1550).

Mackworth, N. H. (1948). The breakdown of vigilance during prolonged visual search. *Quarterly Journal of Experimental Psychology, 1*(1), 6–21.

Mercado, J. E., Rupp, M. A., Chen, J. Y. C., Barnes, M. J., Barber, D., & Procci, K. (2016). Intelligent agent transparency in human-agent teaming for multi-UxV management. *Human Factors, 58*(3), 401–415.

Mishler & Chen, 2023. Effect of automation failure type on trust development in driving automation.

McGarry, S. P. D., Brown, N. L., Drollinger, S. M., Neilson, B. N., Foroughi, C. K., Sibley, C. M., & Coyne, J. T. (2023). The need for specificity in uncertainty communication is further supported: Preliminary evidence from eye tracking. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *67*(1), 453–458. https://doi.org/10.1177/21695067231194332

McGarry, et al., under review

O'Neill, T., McNeese, N., Barron, A., & Schelble, B. (2022). Human-Autonomy Teaming: A Review and Analysis of the Empirical Literature. *Human factors*, *64*(5), 904–938. https://doi.org/10.1177/0018720820960865

Pareek, A., Ro, D. H., Karlsson, J., & Martin, R. K. (2024). Machine learning/artificial intelligence in sports medicine: state of the art and future directions. *Journal of ISAKOS: Joint disorders & orthopaedic sports medicine*, *9*(4), 635–644.

Raven, J., Raven, J. C., & Court, J. H. (1998). Manual for Raven's Progressive Matrices and Vocabulary Scales: Section 4, Advanced Progressive Matrices (1998 ed.). Oxford, UK: Oxford Psychologists Press.

Redick, T. S., Shipstead, Z., Harrison, T. L., Hicks, K. L., Fried, D. E., Hambrick, D. Z., Kane, M. J., & Engle, R. W. (2013). No evidence of intelligence improvement after working memory training: A randomized, placebo-controlled study. Journal of Experimental Psychology: General, 142(2), 359–379.

Redick, T. S., Shipstead, Z., Meier, M. E., Montroy, J. J., Hicks, K. L., Unsworth, N., Kane, M. J., Hambrick, D. Z., & Engle, R. W. (2016). Cognitive predictors of a common multitasking ability: Contributions from working memory, attention control, and fluid intelligence. *Journal of experimental psychology: General*, *145*(11), 1473–1492. https://doi.org/10.1037/xge0000219

Rice, S., & Geels, K. (2010). Using system-wide trust theory to make predictions about dependence on four diagnostic aids. *The Journal of General Psychology, 137*, 362–375.

Sargent, R., Walters, B., & Wickens, C. (2023). Meta-analysis qualifying and quantifying the benefits of automation transparency to enhance models of human performance. In *International Conference on Human-Computer Interaction* (pp. 243–261). Springer Nature Switzerland.

Sibley, C., Coyne, J., & Thomas, J. (2016). Demonstrating the Supervisory Control Operations User Testbed (SCOUT). In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 60, No. 1, pp. 1324–1328).

Thurstone, L. L. (1938). Primary mental abilities. Chicago: University of Chicago Press.

Walliser, J. C., de Visser, E. J., & Shaw, T. H. (2016). Application of a system-wide trust strategy when supervising multiple autonomous agents. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 60, pp. 133–137).

Walliser, J. C., de Visser, E. J., & Shaw, T. H. (2023). Exploring system-wide trust prevalence and mitigation strategies with multiple autonomous agents. *Computers in Human Behavior, 143*, 107671.

Yin, M., Wortman Vaughan, J., & Wallach, H. (2019). Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 chi conference on human factors in computing systems* (pp. 1–12). https://doi.org/10.1145/3290605.3300509