

Automation Reliance Under Time Pressure

Stephen Rice and David Keller
New Mexico State University

A host of factors potentially can affect operator reliance on automation. Some previously studied factors include automation reliability, types of errors, and training. The current study explores another influence on operator reliance—time pressure. Participants performed a simulated target-detection task aided by diagnostic automation that varied in reliability. The amount of time given to operators to make a decision was manipulated to test how time pressure affected operator reliance. The data revealed that when under time pressure, participants tended to depend more on the automation, as seen by increased compliance to the aid's recommendations. This increased compliance benefited overall human-automation performance when the aid was highly reliable but adversely affected overall performance when the aid was less reliable. The data suggest a potential method for mitigating automation under-reliance. Theoretical and applied issues are discussed.

KEYWORDS: Automation, Reliance, Time Pressure

Automation is often defined as an entity that can replace or augment human performance (Wickens & Hollands, 2000). One interesting finding that has emerged in recent studies involving diagnostic automation combined with target detection tasks is the revelation that although overall human-automation performance often exceeds that of the human alone, it does not match, and almost never exceeds, that of the automation itself (e.g., Dixon, Wickens, & McCarley, 2007; Dixon & Wickens, 2006; Rice, in press; Rice, Trafimow, Clayton & Hunt, 2008; Weigmann, McCarley, Kramer, & Wickens, 2006). In other words, in some domains (e.g. luggage screening and aerial reconnaissance), the automation alone is superior to overall performance compared to when a human is paired with the automation and is allowed to override it. This suboptimal performance has raised concerns that perhaps human reliance on automation is flawed; that is, human operators interfere at inappropriate times and second-guess the automation when it is actually correct. In extreme cases, the human may ignore the automation entirely and operate independently of the automation's recommendations (Breznitz, 1984).

Intuitively, one might wonder if perhaps the human should be taken out of the loop altogether. After all, if human-automation performance is inferior to the automation itself, would it not be an improvement to eliminate the human's ability to override the automation? On the face of things, this may seem like an attractive solution; however, there are disadvantages to this proposal. First, when the automation makes a catastrophic error (e.g., the auto-pilot in an aircraft fails),

then it is critical that the human operator be able to take over in this crisis. If the human is locked out of the system, then even obvious errors can result in serious consequences and loss of life. Second, this proposal eliminates all possibility for performance that exceeds that of the automation itself. If the aid has a particular reliability level, then that is the best we can expect of it. Possibly, overall human-automation performance could be improved via practice or training methods, and eventually exceed that of the automation itself. Thus, it seems prudent to analyze potential ways of improving human-automation performance before abandoning the human altogether.

Reliance

There are many variables that can influence operator reliance on automation aids. In Figure 1, we introduce a framework to illustrate a few of the potential factors influencing reliance. One such factor that is commonly studied is trust, which is a mental state. Trust is driven by a subjective assessment of the aid's reliability, and can be defined as an attitude toward automation that includes the belief that the aid will assist one in reaching one's goals (Lee & See, 2004). Intuitively, trust can determine reliance (e.g., high trust leads to high reliance); however, this is not always the case (Weigmann, 2002). Trust itself is also influenced by several variables, including the reliability of the automation itself. For example, Parasuraman, Molloy and Singh (1993) showed that when the reliability of the aid is extremely high, operators can become complacent, such that their over-reliance on the automation can lead to a failure to notice

rare automation errors. Dixon and Wickens (2006) noted that when the reliability of the automation is very low, operators can become under-dependent on the aid, such that they ignore it even when it may be correct.

Another variable that influences trust is the type of error the automation commits. Trust in automation differs as a function of whether the automation commits a false alarm or a miss (Meyer, 2001; 2004; Rice, in press; Parasuraman & Riley, 1997). False alarm prone automation drives down compliance (how operators respond when the automation indicates an event or failure), while miss-prone automation drives down reliance (how operators respond when the automation indicates no event). Rice (in press) demonstrated that the two types of automation errors affected different trust processes, which, in turn, affected different dependent behaviors.

There are other factors that can affect reliance in the absence of a shift in trust. For example, reliance may be unusually high when an operator is engaged in multiple

tasks that require attentional resources to be unequally divided. Imagine a case where a pilot must perform a difficult cognitive task while also maintaining awareness of warning gauges. When the primary task is extremely demanding, pilots may over-depend on the automation during this period of time because they do not have the resources to oversee the automation. Their trust has not shifted but their reliance has.

Other examples could be when the pilot is under immense stress, has higher priorities to attend to, or the consequences of failing to comply are severe. For example, the TCAS (Traffic Collision Avoidance System) is notorious for producing false alarms; however, because of the seriousness of failure to comply (i.e. the two planes crash into each other), pilots must respond to the TCAS alert at all times. An alternative case could be during combat maneuvers. If pilots are concerned about being engaged by an enemy aircraft, they may “let the automation do its thing” while they focus on the more important task of defensive maneuvering.

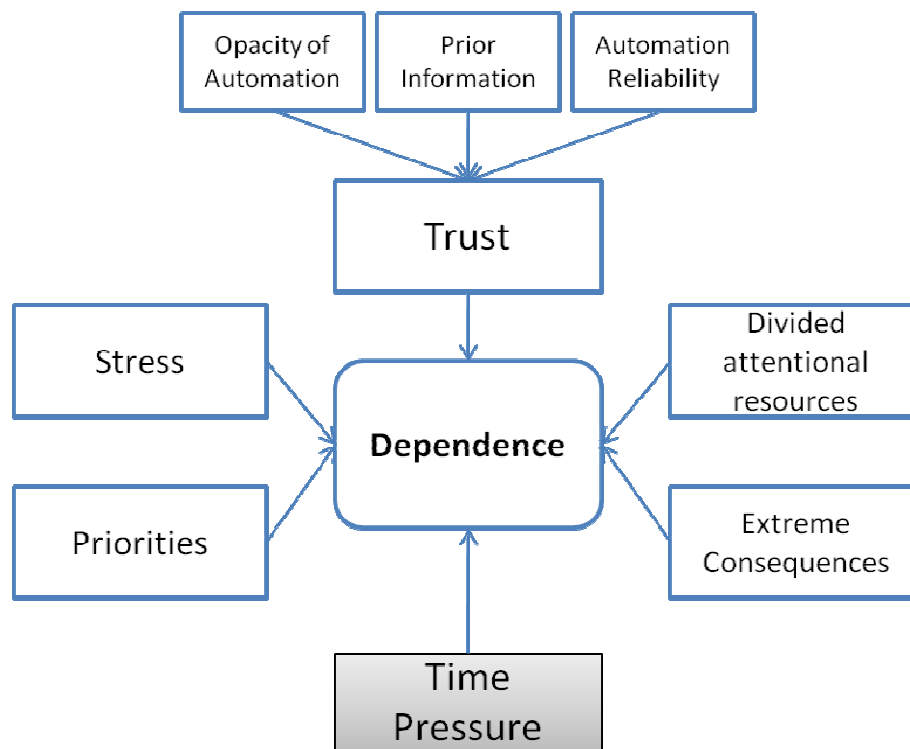


Figure 1. A model of automation reliance.

In summary, the reliance framework that we present is not intended to show all possible determinants of reliance; it merely serves to provide a few examples where trust is a necessary component and few examples where trust does not interact with the outside variables.

Time Pressure

The purpose of this study is to primarily examine one of the factors that we believe is not influenced by trust—time pressure—and thus our initial review focuses on this issue. Research has shown that time pressure has a profound effect on cognitive processes such as decision-making and judgment (Maule, Hockey & Bdzola, 2000; Svenson & Maule, 1993; Maule, 1997; Maule & Edland, 1997). For example, researchers have demonstrated that time pressure can reduce the quality of decision making (Payne, Bettman & Johnson, 1993) and induce less extreme judgments (Kaplan, Wanshula & Zanna, 1993). It has also been shown that the effects of time pressure are not always negative (Maule, et al., 2000). For example Svenson and Benson (1993) found that participants under time-pressure were less likely to engage in framing bias and therefore were able to make better decisions.

Perhaps one reason time pressure affects human performance so much is that it causes them to change the way they process information. Researchers often discuss decision-making and judgment tasks as using two different types of thinking processes (Evans, 2007; Evans, 2008; Evans, Handley & Bacon, 2009; Evans & Over, 1996; Kahneman & Frederick, 2002; Sloman, 1996; Stanovich, 1999). The first type involves slower, less automatic, and sequential thinking called analytic processing. Here, people must make a decision based on information and situations that have not been encountered before. The second type involves fast, unconscious, and automatic thinking, called heuristic processing (Goldstein & Gigerenzer, 2002; Hogarth, 1981).

In the context of this study, the important distinction between the two processes is the time required for each of them. Due to time pressure, the slow analytic process often does not have time to finish (Harreveld, Wagenmakers & Maas, 2007), forcing people to switch strategies to somehow compensate for the increased time pressure (Rieskamp & Hoffrage, 2007). For example, people may alter their decision and judgment process by attempting to gather more information in a shorter amount of time (Ben-Zur & Breznitz, 1981; Edland, 1994; Kerstholt, 1995; Payne, Bettman & Johnson, 1988). People may also attempt to select more important information during the time they have (Ben-Zur &

Breznitz, 1981; Böckenholt & Kroeger, 1993; Kerstholt, 1995; Payne, et al., 1988; Wright, 1974) and rely more on cues that may point to the important information (Maule, 1994; Payne et al., 1988; Rieskamp & Hoffrage, 2007). It is often the case where the human must perform some task (e.g., target detection) in a new environment. In this case a person must turn to the slower analytic (search) process (Harreveld, et al., 2007). Because the analytic search process is slow, it is here that time pressure has the biggest effect (Harreveld, et al., 2007; Ericsson & Staszewski, 1989; Evans, Handley & Bacon, 2009).

When possible, it is often favorable to have a set of accurate heuristics that can be quickly accessed in order to make a decision. A heuristic, as described by Todd and Gigerenzer (2000), is a decision making strategy that “employ[s] a minimum of time, knowledge, and computation to make adaptive choices in real environments” (p. 731). Heuristic reasoning has been posited to reflect a form of rapid, automatic cognitive processing that is distinct from that responsible for more analytic or deliberative judgments (e.g., Evans, 2003), and which dominates decision making when the deliberative system is burdened. Thus, under circumstances in which the heuristic processing system is more reliable than the deliberative system, stressors which hinder deliberative reasoning can, ironically, improve decision making (De Neys, 2006). This implies that if operators can be induced into a heuristic mode of processing, they might be more likely to depend on an automated aid. This should not only increase speed of responding but, when the aid is extremely reliable, increase accuracy as well.

Current Study

In order to test the effects of time pressure on operator reliance, we had participants perform a simulated target-detection task. Participants viewed aerial images of Baghdad, and had to determine whether or not the images contained an enemy target. This task was augmented by a diagnostic aid that provided recommendations as to whether or not a target existed.

We hypothesized that when operators were put under time pressure (in the absence of external stressors or additional tasks), they would comply more often with the aid’s recommendation than their counterparts who were not under time pressure. There was little reason for the operator’s level of trust to change, as the reliability of the aid remained the same across the time manipulation, so any changes in behavior should have been a result of time pressure.

Importantly, this heuristic can be beneficial under certain circumstances, and non-beneficial under other circumstances. By manipulating time and automation reliability, we planned to show how reliance on automation can be profoundly influenced to the benefit *and* decrement of overall operator-automation performance. Thus, our second hypothesis was that, due to the increased reliance, participants under time pressure would perform better than the other participants when the automation was highly reliable, but would perform worse when the reliability of the automation was low.

METHOD

Participants

260 participants (143 females) from a major university in the southwestern United States took part in the experiment for course credit. The mean age was 20.1 ($SD = 2.63$) with a range from 17–38. All participants reported normal or corrected to normal vision.

Design

This study used a 2 x 5 between-participants design, with time-pressure (speeded “s” vs. unspeeded “u”) and automation reliability (100%, 95%, 80%, 65% or no automation “B”) serving as the two independent variables. Participants were randomly assigned to one of the 10 conditions (Bs, Bu, 100s, 100u, 95s, 95u, 80s, 80u, 65s, and 65u). We employed a between-subjects design—instead of a within-subjects design—for two reasons: a) we wanted to avoid contamination effects between the time-pressure conditions (e.g. once participants learn and practice the speeded heuristic, they very well may use it even when the task is unspeeded) and between the reliability conditions, and b) we wanted to keep the experiment down to a reasonable length to avoid issues of fatigue.

Experimental Task

Participants were serially presented with 100 aerial photographs of Baghdad, in which some of the photographs contained an enemy tank. Participants were instructed that if they detected a tank, they should press the “J” key on the provided keyboard. If they determined that no tank was present, they should press the “F” key. Participants were asked to maintain as high accuracy as possible within the time constraints. Participants were also informed that they would be assisted by an automated diagnostic aid that would give them a recommendation on whether or not a tank was present. They were told exactly how reliable the aid was and what types of errors they could expect it to make (false alarms). Lastly, participants were told that they would

either have 2 seconds or 8 seconds to view each image after the automation provided its recommendation.

Once participants read through the instructions, they could press any key to begin the experiment. Each of the 100 trials began with a slide presenting the recommendation of the automation in a text message. It would either state, “The automation has detected a tank!” or “The automation has determined that there is no tank present!” This slide remained for 1000 ms, followed by the aerial photograph.

In the Speeded condition, each photograph remained present for 2 seconds. In the Unspeeded condition, each photograph remained present for 8 seconds. Participants responded by pressing either the “J” key if they detected a tank, or the “F” key if they determined that there was no tank. Following a response, a feedback slide was presented. This slide stated, “Correct!” in green letters or “Incorrect!” in red letters, with the response time for that trial and a running total of percentage correct throughout the experiment. This slide remained present for 1000 ms, followed by the next trial. After 100 trials, the experiment ended automatically.

Apparatus and Stimuli

The experimental display was presented via E-Prime 1.1 on a Dell computer with a 20” monitor, using 1024 x 768 resolution. Accuracy and agreement rates were recorded by E-Prime. Target-absent images were created by using 50 aerial photographs of Baghdad. Target-present images were created by digitally superimposing a simulated tank onto the 50 aerial photographs using Photoshop CS3. Thus, there were 100 photographic stimuli—50 with tanks and 50 without tanks. Figure 2 presents an example image.

Procedures

Participants signed a consent form, and were then seated 21” from the display. Position was controlled with a chin rest. Instructions were given within the e-prime program and all questions were answered verbally by the experimenter. Participants then began the experimental trials. Following the experimental trials, participants were asked to fill out a trust questionnaire (Dixon & Wickens, 2006). They were asked how reliable they believed the automation to be, and how much they trusted the automation. After filling out the questionnaire participants were then debriefed by the experimenter.



Figure 2. A sample image used in the experiment. The tank is located in the top central area to the southeast of the concentric circles.

RESULTS

The following analyses are divided into two parts. First, we discuss accuracy performance in the baseline conditions, followed by similar analyses of the automation conditions. Second, we focus on the reliance effects as seen by the compliance rates when the diagnostic aid provided hits, correct rejections, and false alarms.

Accuracy

Accuracy was measured as the proportion of correct trials to total trials. These data can be found in Table 1 and in Figure 3. The measure d' was not used in this analysis due to the extremely high number of perfect scores in the 100% and 95% conditions. In terms of signal detection theory, d' measures the sensitivity of the participant to distinguish between target-present and target-absent images. Although d' was not used, it should be noted that the effects of d' were almost identical to those of accuracy.

Analysis was first conducted on the Baseline conditions to verify that the task was indeed difficult enough to warrant the assistance of automation and that a traditional speed-accuracy tradeoff (Wickelgren, 1977) would result from the time manipulation. It was clearly the case that when participants were placed under time pressure, their performance suffered tremendously, $t(50) = 4.43$, $p < .001$, $d = 1.25$. In fact, baseline performance in the speeded condition was barely above random performance.

A two-way between-participants ANOVA on the 8 automation conditions, using Reliability and Time as factors, revealed a main effect of Reliability, $F(4, 250) = 161.77$, $p < .001$, $\eta_p^2 = .72$, and no main effect of Time, $F(1, 250) = 1.15$, $p > .10$, $\eta_p^2 = .005$. However, there was an interaction between Reliability and Time, $F(4, 250) = 9.06$, $p < .001$, $\eta_p^2 = .13$, indicating that Time pressure benefited overall performance at some reliability levels, but harmed it at others, an effect clearly seen in Figure 3.

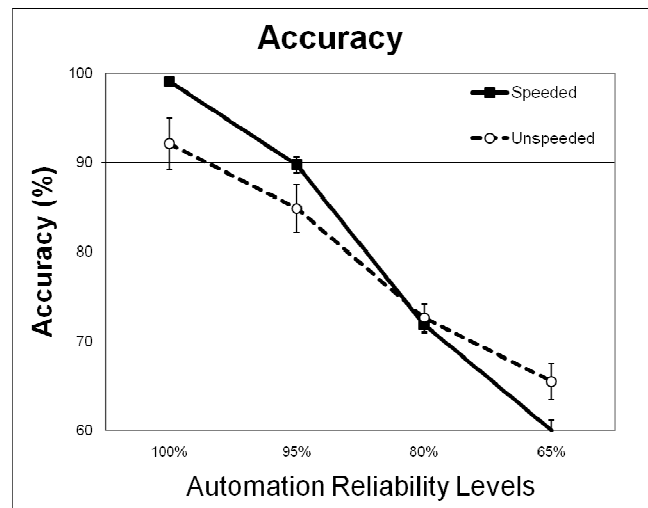


Figure 3. Accuracy data as a function of automation reliability and time. Standard error bars are included.

Planned comparisons revealed that the 100s condition produced higher accuracy than the 100u condition, $t(50) = 2.43$, $p < .01$, $d = .69$, and the 95s condition produced higher accuracy than the 95u condition, $t(50) = 1.73$, $p < .05$, $d = .49$; however, the 65s condition produced lower accuracy than the 65u condition, $t(50) = 2.32$, $p = .01$, $d = .66$. There was no significant difference between the 80s and 80u conditions, $t(50) = .43$, $p > .10$, $d = .12$. The 95s conditions produced higher accuracy than the 80s condition, $t(50) = 13.77$, $p < .001$, $d = 3.89$, which, in turn, produced higher accuracy than the 65s condition, $t(50) = 7.65$, $p < .001$, $d = 2.16$.

Reliance

Operator reliance on the automation was assessed by analyzing participant's agreement rates with the automation. It can be assumed that high reliance on the automation drives high agreement rates (Dixon & Wickens, 2006; Dixon, Wickens & McCarley, 2007). These data can be found in Table 2. It is important to note that when the automation produces hits or correct rejections, then high agreement rates might not necessarily be due to high reliance; that is, it could be that the operator simply performed well in correlation with the performance of the automation. In other words, both the automation and the operator returned the same correct answer independent of each other. Thus, we put special focus on those trials only where the automation produced a false alarm. If the operator agrees with the automation when it fails, then there is no other reasonable explanation for this behavior other than high reliance. Therefore, agreement rates for automation hits, false alarms and correct rejections were analyzed separately.

Table 1. Accuracy as a function of automation reliability level and time pressure. S refers to Speeded, while U refers to Unspeeded

	Reliability Level									
	Baseline-S	Baseline-U	100-S	100-U	95-S	95-U	80-S	80-U	65-S	65-U
Accuracy	51.92	63.54	99.04	92.08	89.73	84.85	71.85	72.62	59.96	65.46
St. Deviation	5.35	12.24	1.15	14.54	4.50	13.68	4.86	7.73	6.26	10.33

Table 2. Agreement rates as a function of automation hits, correct rejections, and false alarms. Note that false alarms were not possible in the 100% reliable conditions. S refers to Speeded, while U refers to Unspeeded

	Reliability Level								
	100-S	100-U	95-S	95-U	80-S	80-U	65-S	65-U	
Hits	0.99	0.93	0.93	0.88	0.82	0.78	0.72	0.74	
Correct Rejections	0.99	0.91	0.94	0.89	0.89	0.88	0.77	0.77	
False Alarms	N/A	N/A	0.95	0.78	0.79	0.64	0.64	0.52	

Automation Hits. Planned comparisons revealed that the 100s condition generated higher agreement rates than the 100u condition, $t(50) = 2.51, p < .01, d = .71$, the 95s condition generated higher agreement rates than the 95u condition, $t(50) = 1.67, p = .05, d = .43$. There were no differences between the 80s and 80u conditions, $t(50) = 1.08, p > .10, d = .31$, or between the 65s and 65u conditions, $t(50) < 1.0$. These data give at least partial evidence that participants under time pressure tend to agree more with the automation.

Automation Correct Rejections. Planned comparisons revealed that the 100s condition generated higher agreement rates than the 100u condition, $t(50) = 2.33, p = .01, d = .66$, the 95s condition generated marginally higher agreement rates than the 95u condition, $t(50) = 1.50, p = .07, d = .42$. There were no differences between the 80s and 80u conditions, $t(50) < 1.0$, or between the 65s and 65u conditions, $t(50) < 1.0$. These data are consistent with the automation hits data.

Automation False Alarms. Because agreement rates during the trials when the automation was correct could occur for reasons other than compliance, as described earlier, we analyzed the automation false alarm trials separately. Planned comparisons revealed that the 95s condition generated higher agreement rates than the 95u condition, $t(50) = 3.04, p < .01, d = .86$, the 80s condition generated higher agreement rates than the 80u condition, $t(50) = 2.51, p < .01, d = .71$, and the 65s condition generated higher agreement rates than the 65u condition, $t(50) = 1.84, p < .05, d = .52$. These data confirm that

when under time pressure, participants tended to agree more with the automation even when it was incorrect.

Trust Questionnaires

Data from the trust questionnaires were straightforward. Because participants were told a priori how reliable the automation was, it was not surprising that their ratings of reliability after the experiment did not differ significantly as a function of time manipulation (all $ps > .10$). Furthermore, participants' ratings of general trust in the automation also did not differ significantly as a function of time manipulation (all $ps > .10$). This indicates that the time manipulation did not shift trust attitudes, but instead directly affected operator reliance in absence of the influence of trust.

DISCUSSION

Theoretical Issues

We first discuss the theoretical implications of the current study, followed by discussion of the practical implications of this study. It is intuitive in many cases that outside factors can affect trust, which, in turn, affects operator reliance on automation. This has been shown multiple times in various paradigms (e.g., Dixon & Wickens, 2006; Dixon, Wickens & McCarley, 2007; Parasuraman, Molloy & Singh, 1993; Parasuraman & Riley, 1997; Rice, in press; Rice, Trafimow, Clayton & Hunt, 2008).

It is less intuitive that these outside factors also can directly affect operator reliance without the influence of

trust. As can be seen in Figure 1, trust plays no role in these tasks, and reliance is driven solely by the outside variable. For example, it easily could be the case that one has no choice but to depend on the automation, despite having little trust in it. Previously, we mentioned example cases that might be driven by divided attentional resources, stress, priorities, and consequences. The current study focuses on another variable—time pressure—that also directly affects operator reliance.

The data support the claim that time pressure directly affected operator reliance, while having no apparent effect on operator perceived trust. When participants were put under time pressure to make quick decisions, they tended to comply more with the automation than when they had more time to search each image for the target. Whereas this compliance during trials when the automation either produced a hit or a correct rejection could be explained by outside factors (independent performance correlations between the human and automation), the fact that operator reliance increased dramatically during trials when the automation *failed* is clear indication of increased reliance on the aid during time pressure.

This increased reliance had some interesting benefits. In the baseline conditions, we see that performance under time pressure suffered greatly relative to performance in the unspeeded condition, an effect due to the traditional speed-accuracy tradeoff (Wickelgren 1977). However, when the automation was highly reliable (100% or 95%), this tradeoff was reversed; that is, overall human-automation performance was superior as a function of less time. Clearly, the reason for this was the increased reliance on the automation, which was presumably superior to the unaided human. By simply putting participants under a time limit, we were able to eliminate the speed-accuracy tradeoff and, in fact, produce significant gains in overall performance relative to unspeeded conditions.

Unfortunately, this increased reliance also had a downside. When the automation was unreliable, participants depended on the automation to the detriment of overall performance. Because the automation was so poor, it did little to help participants, and the added time in the unspeeded condition allowed participants to override the automation and perform better than their counterparts in the speeded condition.

In order to assure that trust did not play a role in the manipulation of operator reliance, we did three things. First, we told the participants a priori exactly how

reliable the automation was and what type of errors to expect. Second, we gave participants trial-by-trial feedback that informed them whether they were answering correctly or not. This allowed them to calibrate to the automation and confirm its reliability level. Third, and most importantly, we asked the participants after the experiment how reliable they believed the automation to be, and how much trust they placed in it. The data were clear -- participants did not think the aids in the speeded conditions were any less reliable or untrustworthy than the aids in the unspeeded conditions.

Practical Issues

From the results of this study, we can identify at least two possible practical implications. First, designers must be aware of how the environment may influence reliance on automation, regardless of how much an operator trusts the automation. In situations where there is time pressure (e.g., emergency situations) we would anticipate that an operator would increase their reliance on an automated system even if the reliability of the automation does not warrant such reliance or when the operator has little trust in the automation. As the results suggest, this may not be a bad thing unless the automation performance is less than what the operators could do by themselves. The opposite scenario also can be imagined. When operators are left in situations with plenty of time, their reliance on the automation may decrease, even if the automation is highly reliable, resulting in less than optimal performance. Designers should take into account the time pressure demands on the operator and adjust accordingly.

Second, in addition to making adjustments to the environment to assure the appropriate level of reliance, the data from this experiment also reveal a potentially beneficial method for training operators to optimize their human-automation performance via a time pressure heuristic. This study suggests that one training technique to increase reliance on automation may be to have operators apply a time-pressure heuristic, especially when the automation is known to be highly reliable. Future research needs to be conducted to determine if this training method would be effective in the long term, and if the benefits of the heuristics training would remain even after the time pressure has been removed. We note that the time pressure heuristic would not be beneficial to overall human-automation performance when the automation is less reliable than the human. Thus, when operators are exposed to less reliable automation, they may need a different training method in order to optimize their compliance with the aid without harming overall performance.

CONCLUSIONS

The results of the experiment show that time pressure does induce reliance on automation, in the absence of any noticeable effects on operator trust. However, this only helped performance when the automation was highly reliable. Performance suffered when the automation was not highly reliable. More obvious ways to increase reliance (e.g., shifting trust or inducing stress on the operator) are usually difficult or impossible to implement, so this method provides a relatively simple way of training operators to depend more heavily on automation when it is necessary and beneficial. From an applied standpoint, the issue of quickly and accurately detecting dangers from aerial views is an important real world problem for those in aviation, intelligence, meteorology, etc. Understanding how operators may overcome the innate tendency to second guess highly reliable systems is essential.

REFERENCES

- Ben Zur, H., & Breznitz, S.J. (1981). The effect of time pressure on risky choice behavior. *Acta Psychologica*, 47, 89-104.
- Böckenholt, U., & Kroeger, K. (1993). The effect of time pressure in multiattribute binary choice tasks. In O. Svenson & A. J. Maule (Eds.), *Time pressure and stress in human judgment and decision making* (pp. 195–214). New York: Plenum Press.
- Breznitz, S. (1983). *Cry-wolf: The psychology of false alarms*. Hillsdale, NJ: Erlbaum.
- De Neys, W. (2006). Dual processing in reasoning – Two systems but one reasoner. *Psychological Science*, 17, 428–433.
- Dixon, S., & Wickens, C. (2006). Automation reliability in unmanned aerial vehicle control: A reliance-compliance model of automation dependence in high workload. *Human Factors*, 48, 3, 474-486.
- Dixon, S., Wickens, C., & McCarley, (2007). On the independence of compliance and reliance: Are automation false alarms worse than misses? *Human Factors*, 49(4), 564-572.
- Edland, A.C. (1994). Time pressure and the application of decision rules. *Scandinavian Journal of Psychology*, 35, 281-291.
- Ericsson, K.A., & Staszewski, J. J. (1989). Skilled memory and expertise: mechanisms of exceptional performance. In D. Klahr, & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon*. Hillsdale: Erlbaum.
- Evans, J. St. B.T. (2003). In two minds: Dual processing accounts of reasoning. *Trends in Cognitive Sciences*, 7, 454-459.
- Evans, J. St. B.T. (2007). *Hypothetical thinking: Dual process in reasoning and judgement*. Hove, UK: Psychology Press.
- Evans, J. St. B.T. (2008). Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology*, 59, 255–278.
- Evans, J. St. B.T., & Over, D.E. (1996). *Rationality and reasoning*. Hove: Psychology Press.
- Evans, J. St. B.T., Handley, S. J., & Bacon, A.,M. (2009). Reasoning under time pressure: A study of causal conditional inference. *Experimental Psychology*, 56(2), 77-83.
- Goldstein, D.G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109, 75-90.
- Harrevel, F., Wagenmakers, E., & van der Maas, H. (2007). The effects of time pressure on chess skill: An investigation into fast and slow processes underlying expert performance. *Psychological Research*, 71, 591-597.
- Hogarth, R. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin*, 90(2), 197-217.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgement. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). Cambridge: Cambridge University Press.
- Kaplan, M. F., Wanshula, L.T., & Zanna, M.P. (1993). Time pressure and information integration in social judgement. In O. Svenson & A. J. Maule, *Time pressure and stress in human judgement and decision making* (pp. 255-267). New York: Plenum.
- Kerstholt, J.H. (1995). Decision making in a dynamic situation: The effect of false alarms and time pressure. *Journal of Behavioral Decision Making*, 8, 181–200.
- Lee, J., & See, K. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46, 50-80.
- Maule, A.J. (1994). A componential investigation of the relation between structural modelling and cognitive accounts of human judgement. *Acta Psychologica*, 87, 199–216.
- Maule, A. J. (1997). Strategies for adapting to time pressure. In R. Flin, E. Salas, M. Strub, & L. Martin, *Decision making under stress: emerging themes and applications* (pp. 271-279). Ashgate: Aldershot.
- Maule, A.J., & Edland, A.C. (1997). The effects of time pressure on judgment and decision making. In R. Ranyard, W. R. Crozier & O. Svenson, *Decision making: cognitive models and explanation* (pp. 189-204). London: Routledge & Kegan Paul.
- Maule, A., Hockey, G., & Bdzola, L. (2000). Effects of time-pressure on decision-making under uncertainty: Changes in affective state and information processing strategy. *Acta psychological*, 104, 283-301.
- Meyer, J. (2001). Effects of warning validity and proximity on responses to warnings. *Human Factors*, 43, 563-572.
- Meyer, J. (2004). Conceptual issues in the study of dynamic hazard warnings. *Human Factors*, 46, 196-204
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39, 230-253.

- Parasuraman, R., Molloy, R., & Singh, I.L. (1993). Performance consequences of automation-induced "complacency." *International Journal of Aviation Psychology*, 3(1), 1-23.
- Payne, J.W., Bettman, J.R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 14, 534-552.
- Payne, J.W., Bettman, J.R., & Johnson, E.J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.
- Rice, S. (in press). Examining single and multiple-process theories of trust in automation. *Journal of General Psychology*.
- Rice, S., Trafimow, D., Clayton, K., & Hunt, G. (2008). Impact of the contrast effect on trust ratings and behavior with automated systems. *Cognitive Technology*, 13(2), 30-41.
- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta psychologica*, 127, 258-276.
- Sloman, S.A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.
- Stanovich, K.E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahway, NJ: Lawrence Erlbaum Associates.
- Svenson, O., & Benson, L. (1993). Framing and time pressure in decision making. In O. Svenson & A. J. Maule, *Time Pressure and Stress in Human Judgment and Decision Making* (pp. 133-144). New York: Plenum.
- Svenson, O., & Maule A.J. (1993). *Time Pressure and Stress in Human Judgment and Decision Making*. New York: Plenum.
- Todd, P.M., & Gigerenzer, G. (2000). Précis of Simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23, 727-780.
- Wickelgren, W.A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, 41, 67-85.
- Wickens, C.D., & Hollands, J. (2000). *Engineering psychology and human performance* (3rd ed.). New York: Prentice Hall.
- Wiegmann, D. (2002). Agreeing with automated diagnostic aids: A study of user's concurrence strategies. *Human Factors*, 44(1), 44-50.
- Wiegmann, D., McCarley, J.S., Kramer, A.F., & Wickens, C.D. (2006). Age and automation interact to influence performance of a simulated luggage screening task. *Aviation, Space, and Environmental Medicine*, 77, 825-831.
- Wright, P.L. (1974). The harassed decision maker: time pressures, distractions and the use of evidence. *Journal of Applied Psychology*, 59, 555-561.

AUTHOR NOTES

The authors wish to thank Amy Wells, Gayle Hunt and Jackie Chavez for their help in collecting data. We also wish to thank David Trafimow and two anonymous reviewers for their helpful comments during the writing of this paper. This research was funded by an Air Force grant (Index #111915). Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors. Direct Correspondence to: Stephen Rice, Ph.D., Department of Psychology, MSC 3452, New Mexico State University, PO Box 30001, Las Cruces, NM 88003-8001. EMAIL: sc_rice@yahoo.com