Influencing Trust for Human–Automation Collaborative Scheduling of Multiple Unmanned Vehicles

Andrew S. Clare, Massachusetts Institute of Technology, Cambridge, **Mary L. Cummings**, Duke University, Durham, North Carolina, and **Nelson P. Repenning**, Massachusetts Institute of Technology, Cambridge

Objective: We examined the impact of priming on operator trust and system performance when supervising a decentralized network of heterogeneous unmanned vehicles (UVs).

Background: Advances in autonomy have enabled a future vision of single-operator control of multiple heterogeneous UVs. Real-time scheduling for multiple UVs in uncertain environments requires the computational ability of optimization algorithms combined with the judgment and adaptability of human supervisors. Because of system and environmental uncertainty, appropriate operator trust will be instrumental to maintain high system performance and prevent cognitive overload.

Method: Three groups of operators experienced different levels of trust priming prior to conducting simulated missions in an existing, multiple-UV simulation environment.

Results: Participants who play computer and video games frequently were found to have a higher propensity to overtrust automation. By priming gamers to lower their initial trust to a more appropriate level, system performance was improved by 10% as compared to gamers who were primed to have higher trust in the automation.

Conclusion: Priming was successful at adjusting the operator's initial and dynamic trust in the automated scheduling algorithm, which had a substantial impact on system performance.

Application: These results have important implications for personnel selection and training for futuristic multi-UV systems under human supervision. Although gamers may bring valuable skills, they may also be potentially prone to automation bias. Priming during training and regular priming throughout missions may be one potential method for overcoming this propensity to overtrust automation.

Keywords: human supervisory control, unmanned vehicles, mixed-initiative planning, priming, gaming

Address correspondence to Andrew S. Clare, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 33-309, Cambridge, MA 02139, USA; e-mail: aclare@alum.mit.edu.

HUMAN FACTORS
Vol. 57, No. 7, November 2015, pp. 1208–1218
DOI: 10.1177/0018720815587803

INTRODUCTION

The use of unmanned vehicles (UVs) has grown dramatically over the past decade. Common uses for UVs will soon include agriculture, wildlife monitoring, firefighting, search and rescue, border patrol, atmospheric research, entertainment, and cargo delivery (Jenkins, 2013). Current UVs typically require multiple human operators, often more than a comparable manned vehicle would require (Haddal & Gertler, 2010).

A potential future method for a single human operator to control multiple heterogeneous (air, land, sea) UVs involves the operator guiding an automated scheduler (AS) in a collaborative process to create, modify, and approve schedules for the team of UVs, which are then carried out by the semiautonomous UVs. Although this concept is known by many names, including "human–automation collaboration" (Miller & Parasuraman, 2003), "human–computer collaboration" (Silverman, 1992), "human guided algorithms" (Klau, Lesh, Marks, & Mitzenmacher, 2003; Thorner, 2007), and "mixed-initiative planning" (Carbonnell & Collins, 1970; Riley, 1989), all such systems involve a human working collaboratively with an optimization algorithm to solve a complex problem or make a decision.

Because of the interaction component, human trust in the AS will be a crucial driver of performance in such futuristic human–automation collaborative systems. Although there are some similarities to the concept of trust between two humans, there are also some significant differences between human–human trust and human–automation trust (Muir, 1987). Human trust in an AS can be defined as the "attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability" (Lee & See, 2004, p. 51). This paper distinguishes trust in a scheduling algorithm for remotely controlling UVs from recent research on human trust in embodied agents (de Visser et al., 2012) or robots (Hancock et al., 2011), as the operator is not colocated with the vehicle and there is no physical embodiment of the algorithm. Operator trust in the AS can fluctuate due both to the operator's initial trust level in the AS and to the behavior of the AS throughout a mission. This phenomenon has been observed in data analysis from a previous human-in-the-loop experiment and has been linked to changes in performance (Gao, Clare, Macbeth, & Cummings, 2013).

Trust calibration is essential in such systems, as either overtrust or undertrust in automation can be detrimental to system performance. Low human trust in the AS can be caused by automation "brittleness," in that the AS can take into account only those quantifiable variables, parameters, objectives, and constraints identified in the design stages that were deemed to be critical (Scott, Lesh, & Klau, 2002; Silverman, 1992; Smith, McCoy, & Layton, 1997). Looking to future scenarios of multi-UV supervision—where unanticipated events, such as weather changes, vehicle failures, unexpected target movements, and new mission objectives, will often occur—the AS may have difficulty accounting for and responding to unforeseen changes in the environment. Operators with low trust may spend an excessive amount of time replanning or adjusting the schedule (Clare, Macbeth, & Cummings, 2012; Cummings, Clare, & Hart, 2010), which can lead to cognitive overload.

Also, overtrust in automation has been cited in a number of costly and deadly accidents in a variety of domains (Cummings, 2004; Parasuraman & Riley, 1997). Overtrust in the AS can lead to the phenomenon of automation bias (Mosier, Skitka, Heers, & Burdick, 1998), where operators disregard or do not search for contradictory information in light of an AS-generated solution that is accepted as correct. A number of empirical studies have shown that when working with imperfect automation, automation bias occurs (Chen & Terrence, 2009; Lee & Moray, 1994; Muir & Moray, 1996; See, 2002).

Given the need to design systems to achieve the "appropriate" level of human trust (Lee & See, 2004), calibrating operators to the reliability of the AS under different situations is essential to achieving high performance in a human–automation collaborative scheduling system. To this end, we investigate the impact of operator trust on human–automation collaboration given a fast but suboptimal AS. The results of a human-in-the-loop experiment are presented in which we evaluated the use of priming to influence operators' initial trust in the AS and the resulting system performance.

TRUST AND PRIMING

An operator's initial trust level in an automated system can vary widely based on the operator's prior knowledge, past experiences, and training (Lee & Moray, 1994; Moray, Inagaki, & Itoh, 2000). Trust is dynamic, however, and can fluctuate throughout a mission based on the operator's perception of how well the AS is performing (Lee & Moray, 1992; Muir & Moray, 1996). A number of studies have shown that human trust has inertia, whereby automation errors do not necessarily cause instantaneous loss in trust, but recovery in trust from severe failures can also be slow (Hoffman, Johnson, Bradshaw, & Underbrink, 2013; Lee & Moray, 1992, 1994; Lewandowsky, Mundy, & Tan, 2000; Parasuraman, 1993; See, 2002).

Priming, which is one way to experimentally influence this initial trust in a system, has been studied extensively in the psychology and neuroscience domains (Cave, 1997; Henson, 2003; Kosslyn & Rosenberg, 2011; Schacter, 1987; Schacter & Buckner, 1998), and it is known that humans are susceptible to anchoring biases in decision making and judgments under uncertainty (Dzindolet, Pierce, Beck, & Dawe, 2002; Tversky & Kahneman, 1974). However, there has been little research on the impact of priming on operators controlling multiple UVs.

In addition, a few studies have involved the impact of framing on human decision making and reliance on an automated decision aid (Dzindolet et al., 2002; Lacson, Wiegmann, & Madhavan, 2005). In these experiments, participants were provided with information about previous automation performance with either positive framing ("the aid usually made about half as many errors as most participants") or negative framing ("the aid usually made about 10 errors in 200 trials"). These studies found that the manner in which information about the reliability of the automation was presented to operators could subtly influence reliance on the automation, but all experiments focused on signal detection rather than the more complex decision making required for controlling multiple UVs.

In another study, Rice, Clayton, Wells, and Keller (2008) primed test participants with images of automation with either positive or negative affect. They found that operators primed with positive images had faster reaction times and higher accuracy in a visual identification task with the assistance of an automated identification aid. However, this was the only task that the operators were conducting, as opposed to the test bed described in this paper where operators were multitasking. Also, the automation was for target identification and had 100% reliability, as opposed to the automated scheduling algorithm used in this test bed, which has been found to be provably good but suboptimal (Choi, Brunet, & How, 2009; Whitten, 2010).

In this experiment, described in more detail in the next section, we elected to use three levels of priming as the main independent variable: "positive priming," "negative priming," and "no priming." The priming verbiage consisted of actual written quotes from operators in a previous experiment using the same test bed (Clare, Cummings, How, Whitten, & Toupet, 2012). For the positive-priming level, the quotes reflected positive naturalistic impressions of the AS, for example, "The system is easy to use and intuitive to work with" and "The automated scheduler was very fast." For the negative-priming level, the quotes reflected dissatisfaction with the AS, for example, "I did not always understand decisions made by the automated scheduler . . . namely, it would not assign tasks . . . while some vehicles were seemingly idle." The no-priming level served as a control group, whereby operators did not receive a passage to read after training. This was a between-participants factor, in that a particular participant experienced only one priming level to avoid training biases and confusion.

THE EXPERIMENT

This section describes an experiment to evaluate the validity of using priming to influence operator trust and system performance in a human–automation collaborative system for controlling multiple UVs.

Apparatus

This study utilized a collaborative, multiple-UV system called Onboard Planning System for UVs Supporting Expeditionary Reconnaissance and Surveillance (OPS-USERS), which leverages decentralized algorithms for vehicle routing and task allocation (Cummings, How, Whitten, & Toupet, 2012). This system functions as a computer simulation but also supports actual flight and ground capabilities (How et al., 2009); all the decision support displays described here have operated actual small air and ground UVs in real time (Kopeikin, Clare, Toupet, How, & Cummings, 2012).

The operator is assisted by an AS in scheduling tasks for the UVs called the consensus-based bundle algorithm, which is a decentralized, polynomial-time, market-based protocol (Choi et al., 2009). A *goal-based* architecture was implemented whereby the human operator guides the high-level goals of the team of UVs (as opposed to guiding each individual vehicle) and the AS assumes the bulk of computation for optimization of task assignments. The AS is responsible for decisions requiring rapid calculations or optimization, and the human operator supervises the AS for high-level goal achievement, including where to search, which tasks get included in the overall plan, and approval of weapons release. More details on the AS can be found in Whitten (2010), with details of the OPS-USERS automation architecture in Cummings et al. (2012).

The primary interface used by the operator is a map display (Figure 1).

Operators had two exclusive tasks that could not be performed by automation: target identification and approval of weapons launch to destroy a hostile target. Operators could also create search tasks, which dictated on the map those areas that the operator wanted the UVs to specifically search. An instant messaging "chat" communication tool provided high-level direction and intelligence to the operator. A primary assumption was that operators had minimal time to interact with the displays due to other mission-related tasks. In order to aid the operator in understanding UV progress toward mission goals, a decision support interface,

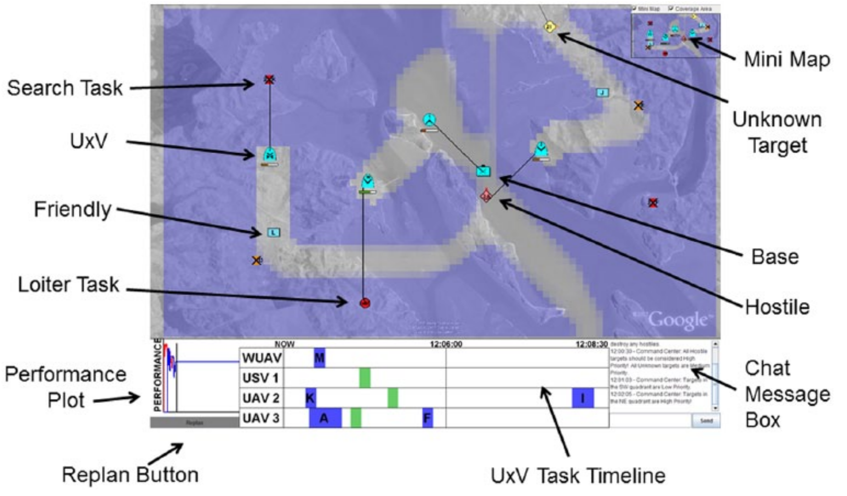

Figure 1. Map display.

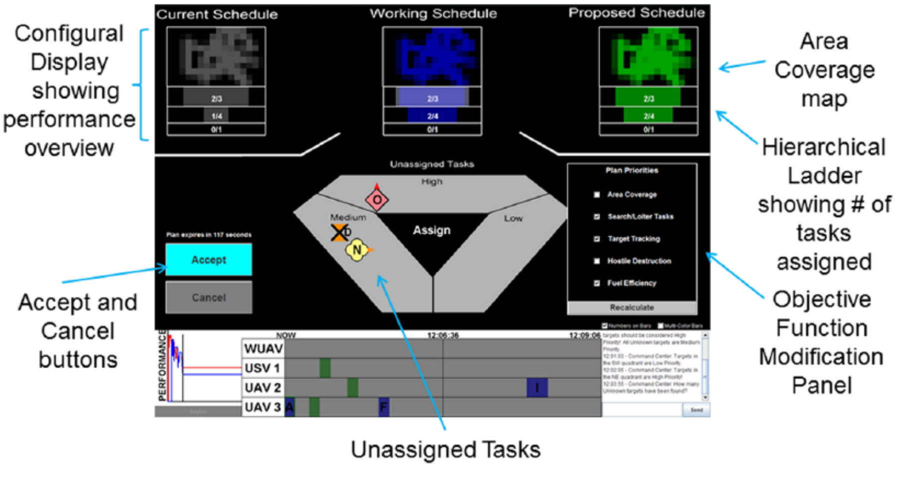

Figure 2. Schedule Comparison Tool (SCT).

called the Schedule Comparison Tool (SCT), was developed (Figure 2). Here the operator could conduct a "what-if" query, which forced the automation to generate a new plan if possible. Operators could also modify the objective function that the AS uses to evaluate schedules for the UVs. Further details of the OPS-USERS interface design and usability testing can be found in previous publications (Clare, Cummings, et al., 2012; Cummings et al., 2010).

Participants

The concept of multiple-UV supervisory control through a decentralized network is a futuristic concept without current subject matter experts. Thus 48 participants were recruited from a northeastern university setting that consisted of 35 men and 13 women, ranging in age from 18 to 32 years with an average age of 23.1 and a standard deviation of 3.8. Each participant filled out a demographic survey prior to the experiment that included age, gender, occupation, military experience, average hours of television viewing, video gaming experience, and perception of UVs.

Experimental Design

Previously, during the development of a computational system dynamics model of

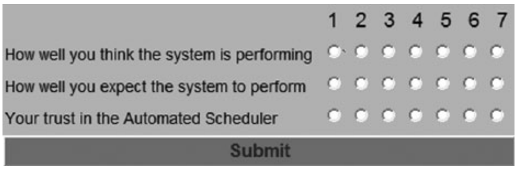

Figure 3. Pop-up survey window.

human–automation collaborative scheduling (Clare, 2013), it was found that factors such as trust, expectations of performance, perceptions of performance, and cognitive workload were all potentially important in determining operator behavior and overall system performance. Thus, the main dependent variables for this experiment were mission performance, primary workload, and subjective ratings, taken both during the missions and postmission.

Overall mission performance was measured through percentage of area coverage, percentage of targets found, percentage of time that targets were tracked, number of correct hostile targets destroyed, and number of mistaken targets destroyed. The primary workload measure was a utilization metric calculating the ratio of the total operator "busy time" to the total mission time. Operators were considered "busy" when performing any task that required direct manual interaction, such as creating search tasks or typing in the chat box. All interface interactions were via a mouse with the exception of the chat messages, which required keyboard input.

Throughout the mission, a pop-up survey window (Figure 3) appeared in the lower left corner of the map display to ask the operator to provide three ratings.

This survey was used to gather near-real-time data on the operator's perception of performance, expectations of how well the system should be performing, and trust in the AS. A Likert rating scale of 1 to 7 (*low* to *high*) was used because that scale is used in an empirically validated and commonly used trust-in-automation questionnaire (Jian, Bisantz, & Drury, 2000). These questions were asked every 2 min, starting at 60 s into the mission. The goal was to sample the operator's perceptions, expectations, and trust level as frequently as possible without distracting the operator from his or her primary tasks. Online probes to gather subjective ratings are commonly used in experiments such as these to measure workload and situation awareness (Endsley, Sollenberger, & Stein, 2000) and have been proposed as a method to measure trust (Miller & Perkins, 2010).

A survey was provided at the end of each mission asking the participant for a subjective rating of his or her confidence, workload, and satisfaction with the plans generated by the AS on a Likert scale from 1 to 5 (*low* to *high*). At the end of the experiment, participants filled out a 12-question survey that is commonly used to measure trust in automation and has been empirically validated (Jian et al., 2000).

In order to familiarize each participant with the interface, a self-paced, slide-based tutorial was provided. Participants then conducted a 15-min practice session during which the experimenter walked the participant through all the necessary functions. Each participant was given the opportunity to ask the experimenter questions during the tutorial and practice session. Each participant also had to pass a proficiency test, which was a five-question slide-based test.

The actual experiment for each participant consisted of two 20-min sessions. During testing, the participant was not able to ask the experimenter questions about the interface and mission. All data and operator actions were recorded, and Camtasia© was used to record the operator's actions on the screen. Participants were paid $10 per hour for the experiment, and a performance bonus of a $100 gift card was given to the individual who obtained the highest mission performance metrics (to encourage maximum effort).

Trust Priming Results

As expected, results showed that participants who experienced the positive-priming level had higher ratings of trust in the AS just prior to the actual experiment. Pairwise Mann-Whitney comparisons showed that the positive-priming group had 13.8% higher trust ratings ($M = 6.19$, $SD = 0.66$) compared to the no-priming control group ($M = 5.44$, $SD = 0.81$; $Z = -2.570$, $p = .010$) and 20.7% higher trust ratings compared to the negative priming group ($M = 5.13$, $SD = 0.92$; $Z = -3.186$, $p = .002$. There were no significant differences in pre-experiment trust ratings between the negative-priming group and the no-priming control group ($Z = -0.807$, $p = .420$).

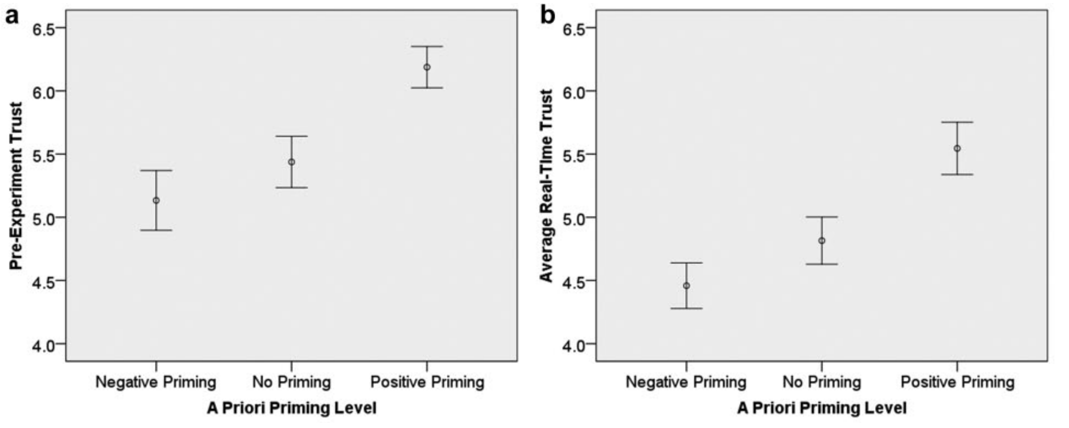

Figure 4. Trust ratings comparison. (a) Pre-experiment ratings. (b) Average real-time ratings during the mission. Standard error bars are shown. $N = 32$ for each a priori priming level.

More importantly, similar results were found for the average real-time rating of trust during the missions, whereby the positive-priming group had 14.9% higher trust ratings ($M = 5.54$, $SD = 1.17$) compared to the no-priming control group ($M = 4.82$, $SD = 1.04$; $Z = -2.614$, $p = .009$) and 24.2% higher trust ratings compared to the negative-priming group ($M = 4.46$, $SD = 1.01$; $Z = -3.741$, $p < .001$). Again, there were no significant differences in average real-time trust ratings between the negative-priming group and the no-priming control group ($Z = -1.036$, $p = .300$). These results are shown in Figure 4.

It appears that positive a priori priming had the desired effect of raising initial trust in the automation. This higher trust level among the positive-priming group was maintained on average throughout the experiment, although trust across all groups during the mission was lower than pre-experiment trust, as shown in Figure 4b.

Overall, these results show that priming can be effective at influencing initial trust level, especially when the priming is meant to raise trust. Interestingly, after the end of the experiment, there were no significant differences in trust across the three a priori priming groups according to a Kruskal-Wallis omnibus test of the 12-question postexperiment trust survey data, $\chi^2(2, N = 48) = 1.986$, $p = .371$. Thus, these data also provide evidence that operators adjust their trust level over time as they work with the AS, and thus the effects of priming are not enduring. This finding aligns with previous empirical evidence demonstrating that trust has inertia (Lee & Moray, 1994; Lewandowsky et al., 2000; Parasuraman, 1993), whereby the effect of perceived automation performance on trust is not instantaneous but changes over time.

System Performance Results

There were no significant differences in the system performance metrics, including the primary metric of area coverage, across the priming groups. However, a post-hoc analysis of the impact of frequency of computer and video game playing provided different results. This analysis was conducted for two main reasons. First, gaming frequency was the most significant demographic predictor of utilization, implying that frequent gamers experienced lower workload levels ($\rho = -0.450$, $p < .001$). Second, in terms of the real-time subjective ratings measured via the interface shown in Figure 3, gaming frequency correlated with higher average ratings of trust ($\rho = 0.216$, $p = .037$), perceived performance ($\rho = 0.248$, $p = .016$), and expectations ($\rho = 0.302$, $p = .003$).

Although all of these correlations were weak to moderate in strength, the fact that gamers were faster at using the interface and had a somewhat higher propensity to trust the automation raised the question of whether frequent gamers reacted to a priori priming differently than nongamers. To facilitate this analysis, participants were divided into categories of gamers and nongamers based on

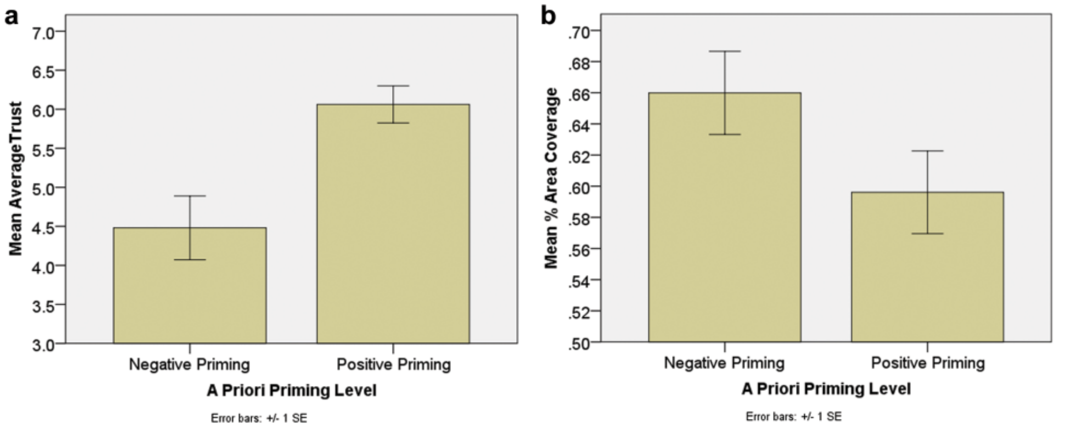

Figure 5. Impact of a priori priming on gamers. (a) Average real-time rating of trust in automated scheduler. (b) Area coverage system performance by the end of the mission. $n = 10$ for negative-priming gamers and $n = 12$ for positive-priming gamers.

participants' self-reported frequency of playing computer and video games. Eighteen participants who reported that they were "weekly gamers," "a-few-times-a-week gamers," or "daily gamers" were classified as gamers. The other 30 participants reported that they played computer or video games once a month or less frequently.

There was a fairly equal distribution of gamer participants across the a priori priming levels: negative priming (5), no priming (7), and positive priming (6). Gamers who experienced the positive-priming level had higher ratings of trust in the AS prior to the experiment. Mann-Whitney pairwise comparisons showed that the positive-priming gamers had 22.8% higher trust ratings ($M = 6.67$, $SD = 0.49$) compared to the no-priming gamers ($M = 5.43$, $SD = 0.76$; $Z = -2.598$, $p = .009$) and 19.1% higher trust ratings compared to the negative-priming gamers ($M = 5.60$, $SD = 0.52$; $Z = -2.364$, $p = .018$). Similar results were found for the average real-time rating of trust during the missions, whereby the positive-priming gamers had 23.9% higher trust ratings ($M = 6.06$, $SD = 0.82$) compared to the no-priming gamers ($M = 4.89$, $SD = 1.21$; $Z = -2.446$, $p = .014$) and 35.2% higher trust ratings compared to the negative-priming gamers ($M = 4.48$, $SD = 1.22$; $Z = -2.882$, $p = .002$). Again, there were no significant differences in average real-time trust ratings between the negative-priming gamers and the no-priming gamers ($Z = -0.536$, $p = .592$).

In terms of system performance, gamers who experienced negative priming had 10% higher average area coverage performance ($M = 66.0\%$, $SD = 8.4\%$) as compared to gamers with positive a priori priming ($M = 59.6\%$, $SD = 9.2\%$), which was a marginally significant difference (Mann-Whitney $Z = -1.715$, $p = .086$). So gamers with positive priming had 35% higher average trust but 10% lower average area coverage performance as compared to gamers with negative priming. These results are shown in Figure 5.

There were no other significant differences in system performance metrics among gamers based on priming. Also, a similar analysis of nongamers revealed that while a priori priming did affect their reported trust in the AS, there were no significant differences in behavior or system performance across the a priori priming groups for nongamers.

Gamers in the negative-priming group understood the imperfections in the automation, reporting lower trust in the AS. However, these differences in performance did not manifest in the nongaming population. Why did priming of trust influence only the behavior of gamers in a way that affected system performance? One potential reason is that gamers may have a higher propensity to overtrust automation. As described earlier, gaming frequency correlated with higher average ratings of trust in the AS ($\rho = 0.216$, $p = .037$). Also, gamers began the mission with significantly higher ratings of trust

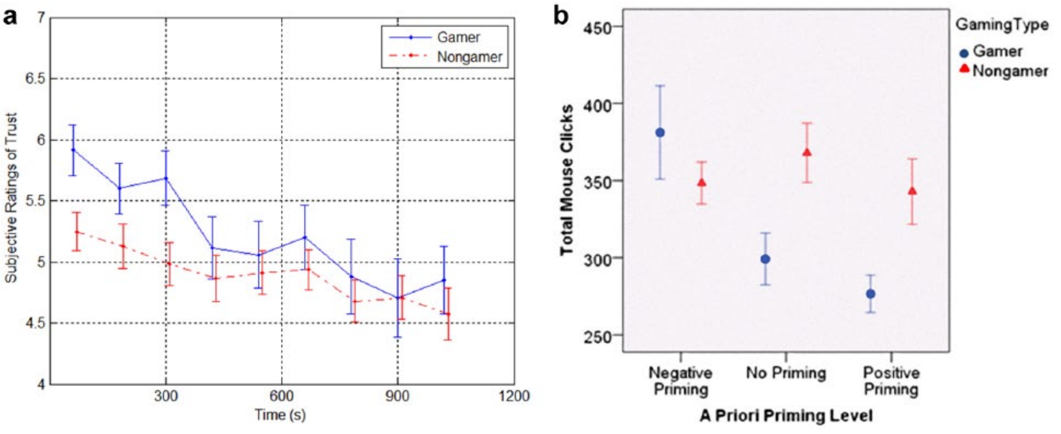

Figure 6. Comparison of gamers and nongamers. (a) Real-time trust ratings. (b) Total mouse clicks. Standard error bars are shown.

(M = 5.91, SD = 1.22) as compared to all nongamers (M = 5.25, SD = 1.16) according to a Mann-Whitney test ($Z = -2.254$, $p = .024$) but eventually adjusted their trust to the same, lower levels of trust of nongamers (Figure 6a).

Additional evidence of the different reactions of gamers and nongamers to priming is provided by an analysis of total mouse clicks. Mouse clicks provide a measure of the level of involvement of the operator and can serve as a proxy for actions to guide the automation (Bruni, Marquez, Brzezinski, & Cummings, 2006; Janzen & Vicente, 1998). An ANOVA indicated a significant difference in the total mouse clicks among the priming levels, $F(2, 90) = 3.765$, $p = .027$. There was also a significant effect for gamer versus nongamer, $F(1, 90) = 4.458$, $p = .038$, and a significant interaction effect between gamer/nongamer and priming level, $F(2, 90) = 4.135$, $p = .019$. Gamers (M = 314, SD = 78.9) overall had 11% fewer total mouse clicks on average as compared to nongamers (M = 352, SD = 79.7). There was no difference in the total mouse clicks of nongamers across the priming levels. However, gamers reacted to negative priming with a significantly larger number of mouse clicks (M = 381, SD = 95.6), whereas gamers with positive priming had the smallest number of mouse clicks (M = 277, SD = 42.1) of any group (Figure 6b). Gamers were prone to overtrusting the automation, and with positive priming, they did not take as much action to guide the automation.

As described previously, overtrust in imperfect automation can lead to the phenomenon of automation bias (Mosier et al., 1998), whereby operators accept an AS-generated solution while disregarding or neglecting to search for contradictory information (Chen & Terrence, 2009; Cummings, 2004; Lee & Moray, 1994; Muir & Moray, 1996; See, 2002). Thus, positive priming may have induced automation bias in gamers, whereby they overtrusted the imperfect AS and did not intervene frequently enough. Negative priming may have pushed gamers to a more appropriate level of trust in the AS, helping them avoid automation bias and encouraging them to intervene more frequently.

Another possible reason is that priming triggered a learned behavior in gamers. Wickens and Hollands (2000) proposed in their human information-processing model that working memory has a more direct impact on perception and response selection as compared to long-term memory. Previous studies have shown that priming can "spread activation," meaning that the prime activates an association in memory prior to carrying out a task (Anderson, 1983; Niedeggen & Rösler, 1999). Additionally, the strength of the effect of the prime on behavior is influenced by the match between earlier experiences and the current situation (Domke, Shah, & Wackman, 1998; Lorch, 1982). Gamers, especially those who play action video games, have learned how to manually control characters or vehicles and manipulate automation to obtain a desired result. Thus, negative priming of gamers may have activated this previously learned behavior to intervene more frequently in order to manipulate the automation to improve performance. In contrast, nongamers likely did not have as much previous experience working with automation, and thus their behavior did not change significantly even though they reported lower trust.

The limitations of this analysis must be considered. There is a need for further experiments, as some priming results have been challenging to replicate in the psychology literature (Kosslyn & Rosenberg, 2011). The definition of a "gamer" is based on self-reported information, meriting further research to establish which types of video games and what frequency of video game play influence operator trust, behavior, and performance. The sample size of gamers in this data set was small, only 18 participants conducting two missions each. Authors of future research should aim to evaluate these findings with a larger sample size of gamers and nongamers. Despite these limitations, initial evidence shows that previous experiences with automation and video game play can have a significant impact on initial trust level in automation and reaction to priming/training methods.

CONCLUSIONS

In this paper we described an experiment that was conducted to evaluate the validity of using priming to influence operator trust and system performance in a human–automation collaborative system for controlling multiple UVs. It was found that priming the initial trust level of operators using quotes from previous users of the system was successful at adjusting initial trust levels, with the strongest effect from positive priming. Participants who play computer and video games frequently were found to have a higher propensity to overtrust automation but also experienced lower workload levels. By priming these gamers to lower their initial trust to a more appropriate level, system performance was improved by 10% as compared to that of gamers who were primed to have higher trust in the AS. Gamers who experienced negative priming took more action to guide the fast, but suboptimal, automation as compared to gamers who experienced positive priming.

These results have interesting implications for personnel selection and training for future real-time human–automation scheduling systems for multiple UVs. Although gamers may bring valuable skills, such as faster visual attentional processing (Green & Bavelier, 2003) and faster encoding of visual information into short-term memory (Wilms, Petersen, & Vangkilde, 2013), they are also potentially prone to automation bias. One potential method for overcoming this propensity to overtrust automation is through priming during training. However, results in this experiment demonstrated that the effects of priming are not enduring; thus regular priming throughout missions may be necessary to maintain the appropriate level of trust, as Rice et al. (2008) proposed.

ACKNOWLEDGMENTS

This research is sponsored by the Office of Naval Research, the Air Force Office of Scientific Research, and Aurora Flight Sciences. Olivier Toupet, formerly of Aurora Flight Sciences, provided extensive algorithm support. Professor John How and the MIT Aerospace Controls Laboratory provided test bed support. The reviewers were also instrumental in shaping this work.

KEY POINTS

- This study presents an experiment in which operators experienced different levels of trust priming prior to conducting simulated missions of supervising a decentralized network of heterogeneous unmanned vehicles with the assistance of an automated scheduling algorithm.
- The results of the experiment showed that participants who play computer and video games frequently had a higher propensity to overtrust automation, and with positive priming, they did not take as much action to guide the automation.
- By priming gamers to lower their initial trust to a more appropriate level, system performance was improved by 10% as compared to that of gamers who were primed to have higher trust in the automation.

REFERENCES

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*, 261–295.

Bruni, S., Marquez, J., Brzezinski, A., & Cummings, M. L. (2006). Visualizing operators' cognitive strategies in multivariate optimization. In *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting* (pp. 1180–1184). Santa Monica, CA: Human Factors and Ergonomics Society.

Carbonnell, J. R., & Collins, A. M. (1970). *Mixed-initiative systems for training and decision-aid applications* (Tech. Rep. ESD-TR-70-373). Bedford, MA: Electronics Systems Division, Laurence G. Hanscom Field, U.S. Air Force.

Cave, C. B. (1997). Very long-lasting priming in picture naming. *Psychological Science*, *8*, 322–325.

Chen, J. Y. C., & Terrence, P. I. (2009). Effects of imperfect automation on concurrent performance of military and robotics tasks in a simulated multi-tasking environment. *Ergonomics*, *52*, 907–920.

Choi, H.-L., Brunet, L., & How, J. P. (2009). Consensus-based decentralized auctions for robust task allocation. *IEEE Transactions on Robotics*, *25*, 912–926.

Clare, A. S. (2013). *Modeling real-time human-automation collaborative scheduling of unmanned vehicles* (PhD thesis). Massachusetts Institute of Technology, Cambridge.

Clare, A. S., Cummings, M. L., How, J., Whitten, A., & Toupet, O. (2012). Operator objective function guidance for a real-time unmanned vehicle scheduling algorithm. *AIAA Journal of Aerospace Computing, Information and Communication*, *9*, 161–173. doi:10.2514/1.I010019

Clare, A. S., Macbeth, J. C., & Cummings, M. L. (2012, June). *Mixed-initiative strategies for real-time scheduling of multiple unmanned vehicles*. Paper presented at the American Control Conference, Montreal, Canada.

Cummings, M. L. (2004, September). *Automation bias in intelligent time critical decision support systems*. Paper presented at the AIAA 3rd Intelligent Systems Conference, Chicago, IL.

Cummings, M. L., Clare, A. S., & Hart, C. S. (2010). The role of human-automation consensus in multiple unmanned vehicle scheduling. *Human Factors*, *52*(1), 17-27.

Cummings, M. L., How, J., Whitten, A., & Toupet, O. (2012). The impact of human-automation collaboration in decentralized multiple unmanned vehicle control. *Proceedings of the IEEE*, *100*, 660–671.

de Visser, E. J., Krueger, F., McKnight, P., Scheid, S., Smith, M., Chalk, S., & Parasuraman, R. (2012). The world is not enough: Trust in cognitive agents. In *Proceedings of the Human Factors and Ergonomics Society 56th Annual Meeting* (pp. 263–267). Santa Monica, CA: Human Factors and Ergonomics Society.

Domke, D., Shah, D. V., & Wackman, D. B. (1998). Media priming effects: Accessibility, association, and activation. *International Journal of Public Opinion Research*, *10*, 51–74.

Dzindolet, M. T., Pierce, L. G., Beck, H. P., & Dawe, L. A. (2002). The perceived utility of human and automated aids in a visual detection task. *Human Factors*, *44*, 79–94.

Endsley, M. R., Sollenberger, R., & Stein, E. (2000, October). *Situation awareness: A comparison of measures*. Paper presented at the Human Performance, Situation Awareness and Automation: User Centered Design for the New Millennium Conference, Savannah, GA.

Gao, F., Clare, A. S., Macbeth, J. C., & Cummings, M. L. (2013, March). *Modeling the impact of operator trust on performance in multiple robot control*. Paper presented at the AAAI Spring Symposium: Trust and Autonomous Systems, Stanford, CA.

Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature*, *423*, 534–537.

Haddal, C. C., & Gertler, J. (2010). *Homeland Security: Unmanned aerial vehicles and border surveillance*. Washington, DC: Congressional Research Service.

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*, 517–527.

Henson, R. N. A. (2003). Neuroimaging studies of priming. *Progress in Neurobiology*, *70*, 53–81.

Hoffman, R. R., Johnson, M., Bradshaw, J. M., & Underbrink, A. (2013). Trust in automation. *IEEE Intelligent Systems*, *28*, 84–88.

How, J. P., Fraser, C., Kulling, K. C., Bertuccelli, L. F., Toupet, O., Brunet, L., Bachrach, A., & Roy, N. (2009). Increasing autonomy of UAVs. *IEEE Robotics & Automation Magazine*, *16*(2), 43–51.

Janzen, M. E., & Vicente, K. J. (1998). Attention allocation within the abstraction hierarchy. *International Journal of Human–Computer Studies*, *48*, 521–545.

Jenkins, D. (2013). *The economic impact of unmanned aircraft systems integration in the United States*. Arlington, VA: Association for Unmanned Vehicle Systems International.

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

Klau, G. W., Lesh, N., Marks, J., & Mitzenmacher, M. (2003). *Human-guided search: Survey and recent results*. Cambridge, MA: Mitsubishi Electric Research Laboratories.

Kopeikin, A., Clare, A., Toupet, O., How, J. P., & Cummings, M. L. (2012, August). *Flight testing a heterogeneous multi-UAV system with human supervision*. Paper presented at the AIAA Guidance, Navigation, and Control Conference, Minneapolis, MN.

Kosslyn, S. M., & Rosenberg, R. S. (2011). *Introducing psychology: Brain, person, group*. Boston, MA: Pearson Learning Solutions.

Lacson, F. C., Wiegmann, D. A., & Madhavan, P. (2005). Effects of attribute and goal framing on automation reliance and compliance. In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting* (pp. 482–486). Santa Monica, CA: Human Factors and Ergonomics Society.

Lee, J. D., & Moray, N. (1992). Trust, control strategies and allocation of functions in human–machine systems. *Ergonomics*, *35*, 1234–1270.

Lee, J. D., & Moray, N. (1994). Trust, self-confidence, and operators' adaptation to automation. *International Journal of Human–Computer Studies*, *40*, 153–184.

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

Lewandowsky, S., Mundy, M., & Tan, G. P. (2000). The dynamics of trust: Comparing humans to automation. *Journal of Experimental Psychology: Applied*, *6*, 104–123.

Lorch, R. F. (1982). Priming and search processes in semantic memory: A test of three models of spreading activation. *Journal of Verbal Learning and Verbal Behavior*, *21*, 468–492.

Miller, C. A., & Parasuraman, R. (2003). Beyond levels of automation: An architecture for more flexible human–automation collaboration. In *Proceedings of the Human Factors and Ergonomics Society 47th Annual Meeting* (pp. 182–186). Santa Monica, CA: Human Factors and Ergonomics Society.

Miller, J. E., & Perkins, L. (2010, June). *Development of metrics for trust in automation*. Paper presented at the International Command and Control Research and Technology Symposium (ICCRTS), Santa Monica, CA.

Moray, N., Inagaki, T., & Itoh, M. (2000). Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied*, *6*, 44–58.

Mosier, K., Skitka, L. J., Heers, S., & Burdick, M. D. (1998). Automation bias: Decision making and performance in high-tech cockpits. *International Journal of Aviation Psychology*, *8*, 47–63.

Muir, B. M. (1987). Trust between humans and machines, and the design of decision aids. *International Journal of Man–Machine Studies*, *27*, 527–539.

Muir, B. M., & Moray, N. (1996). Trust in automation: Part II. Experimental studies of trust and human intervention in a process control simulation. *Ergonomics*, *39*, 429–460.

Niedeggen, M., & Rösler, F. (1999). N400 effects reflect activation spread during retrieval of arithmetic facts. *Psychological Science*, *10*, 271–276.

Parasuraman, R. (1993). Effects of adaptive function allocation on human performance. In D. J. Garland & J. A. Wise (Eds.), *Human factors and advanced aviation technologies* (pp. 147–157). Daytona Beach, FL: Embry-Riddle Aeronautical University Press.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, *39*, 230–253.

Rice, S., Clayton, K., Wells, A., & Keller, D. (2008, May). *Manipulating trust behaviors in a combat identification task*. Paper presented at the Combat Identification Workshop, Gold Canyon, AZ.

Riley, V. (1989). A general model of mixed-initiative human–machine systems. In *Proceedings of the Human Factors and Ergonomics Society 33rd Annual Meeting* (pp. 124–128). Santa Monica, CA: Human Factors and Ergonomics Society.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 501–518.

Schacter, D. L., & Buckner, R. L. (1998). Priming and the brain. *Neuron*, *20*, 185–195.

Scott, S. D., Lesh, N., & Klau, G. W. (2002, April). *Investigating human–computer optimization*. Paper presented at the Conference on Human Factors in Computer Systems, Minneapolis, MN.

See, K. A. (2002). *The Effect of sound on performance, reliance, decision making, and trust in semi-automatic process control systems* (Master of science thesis). University of Iowa, Iowa City.

Silverman, B. G. (1992). Human–computer collaboration. *Human–Computer Interaction*, *7*(2), 165–196.

Smith, P., McCoy, E., & Layton, C. (1997). Brittleness in the design of cooperative problem-solving systems: The effects on user performance. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, *27*, 360–371.

Thorner, J. L. (2007). *Trust-based design of human-guided algorithms* (Master of science thesis). Massachusetts Institute of Technology, Cambridge.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.

Whitten, A. K. (2010). *Decentralized planning for autonomous agents cooperating in complex missions* (Master of science thesis). Massachusetts Institute of Technology, Cambridge.

Wickens, C. D., & Hollands, J. G. (2000). *Engineering psychology and human performance* (3rd ed.). Upper Saddle River, NJ: Prentice Hall.

Wilms, I. L., Petersen, A., & Vangkilde, S. (2013). Intensive video gaming improves encoding speed to visual short-term memory in young male adults. *Acta Psychologica*, *142*, 108–118.

Andrew S. Clare received his PhD in aeronautics and astronautics from the Massachusetts Institute of Technology in 2013. He is currently a management and operations consultant with McKinsey and Company, based in San Francisco.

Mary L. Cummings received her PhD in systems engineering from the University of Virginia in 2004. She is currently an associate professor in the Duke University Department of Mechanical Engineering and Materials Science, the Duke Institute of Brain Sciences, and the Duke Electrical and Computer Engineering Department. She is the director of the Duke Humans and Autonomy Laboratory.

Nelson P. Repenning received his PhD in operations management and system dynamics from the Massachusetts Institute of Technology in 1996. He is currently a professor of system dynamics and an associate professor of management science and organization studies at the MIT Sloan School of Management.

Date received: June 2, 2014
Date accepted: April 3, 2015