

Manipulating Trust Behaviors in a Combat Identification Task

Stephen Rice, Krisstal D. Clayton, Amy Wells, & David Keller

New Mexico State University

ABSTRACT

Objective: To examine how positive and negative images can be used to prime operator trust behaviors towards automation. **Background:** Prior research shows that trust in automation can be affected by various manipulations of the automation, but little has been done in priming participants prior to interactions with automation. **Methods:** Participants were first shown positive or negative images of real-world automation, and then were asked to complete a combat identification task. **Results:** The data revealed that priming positive images resulted in lower trust in the automation, while priming negative images resulted in higher trust. **Conclusions:** Priming operators prior to interactions with automation can have surprising contrast effects depending on the primes and the conditions.

INTRODUCTION

The literature on trust in and dependence on automation has expanded in recent years; however, most of the research has focused on manipulating various properties of the automation (e.g. reliability, bias, etc.) during a human-automation task (Dixon & Wickens, 2006; Dixon, Wickens & McCarley, 2007; Parasuraman & Riley, 1997; Wickens & Dixon, 2007). The current study moves away from this paradigm in order to examine how priming operators prior to the human-automation interaction can also have strong and counter-intuitive effects on overall performance. In particular, we believe it is possible to use simple primes, such as the presentation of positive or negative images, to manipulate operator trust and dependence in contrasting directions.

Contrast effects. According to experimental psychologists, judgments are made in reference to their surrounding context (Stapel & Winkielman, 1998). Some contexts will make a target appear smaller, brighter, faster, or older (Aarts & Dijksterhuis, 2002; Stapel & Winkielman, 1998). For this comparison to occur, the context must be in the same category as the target (Brown, 1953). Therefore, if an architect wanted his or her building to have the illusion of being larger than it is, the architect might build on land that resides between two buildings of smaller stature. Putting the building next to a small tree would not really assist the illusion since the tree is not in the same category as “building”. However, physical properties are not the only aspect influenced by context. Social psychologists have demonstrated that personality characteristics (Stapel & Koomen, 1997), self-esteem (Jones & Buckingham, 2005), performance ratings (Sumer & Knight, 1996), perceptions about in group and out group members (Hicklin & Wedell, 2007), and affect (Geers & Lassiter, 2005) can all be influenced by context.

Interestingly, the context can lead to a target judgment in the same direction (assimilation) or in the opposite direction of the context (contrast).

Contrast effects have been extensively studied by social psychologists (e.g. Becker & Miller, 2002; Brewer & Chapman, 2003; Tormala & Clarkson, 2007). A contrast effect occurs when a person is primed with a stimulus, asked to make a judgment about a target, and judges the target in the opposite direction of the stimulus. For example, Herr, Sherman and Fazio (1983) found that when participants were presented with names of ferocious animals (the prime), they tended to rate other animals (targets) as being less ferocious, and vice versa. This also occurs when people are judging other people, such as in the Tennis and Dabbs (1975) study where females who were primed with unattractive faces showed a contrast effect.

Finally, contrast effects also occur when people are judging their past. For example, much research has focused on what is known as *counterfactual thinking*, or having thoughts that are “literally, contrary to the facts” (Roese, 1997, p. 133). Counterfactual thinking is brought on by negative affect. So, if a person has a bad experience, that person may have an overwhelming need to rid their minds of that negative affect and thereby, make a positive spin out of the situation. Someone who was in a bad car accident, for instance, may focus on the idea that it could have been much worse, or tell themselves and others “I am so lucky that I am alive”. This type of positive spin is referred to as downward counterfactual thinking. However, some people tend to think of how the situation could have been improved, which is referred to as upward counterfactual thinking. Contrary to the downward thinker who was in the car accident, the upward counterfactual thinker might have thoughts such as “If only I had not turned down the wrong street, this never would have happened.” Regardless of downward or upward counterfactual thought, these types of thoughts have a contrast effect in that thinking about what

might have gone worse makes one feel better and thinking about what might have been better makes one feel worse (Walchli & Landman, 2003).

Although these contrast effects on judgments and past experiences may seem trivial, their influence on behavior is impressive (Dijksterhuis, Spears, & Lépinasse, 2001; Dijksterhuis, Spears, Postmes, Stapel, Koomen, van Knippenberg, & Scheepers, 1998). Counterfactual thinking has the capability of assisting people with developing insights into future behavior (Boninger, Gleicher, & Strathman, 1994; Landman, Vandewater, Stewart, & Malley, 1995). For example, if a person experiences a negative situation due to their own devices, counterfactual thinking could inspire that person to think of what they could have done in order to have a positive outcome. Take for example the employee that gets fired instead of promoted due to tardiness. If this employee practices counterfactual thinking, such as “If only I had showed up to work on time, I would have a promotion”, then the next time he/she has a desire to be late for work, negative affective feelings of “what might have been” could possibly influence that person to be on time.

However, contrast effects do not only influence possible future behavior. Dijksterhuis et al. (1998) demonstrated that priming participants with Einstein caused a decrease in immediate performance on an intellectual task. In the same study, walking speed was increased when participants were primed with the elderly Dutch Queen Mother (Dijksterhuis et al., 1998). Schubert and Häfner (2003) found similar results; their participants who were primed with an elderly exemplar had faster reaction times to a word task than participants who were primed with a young exemplar. Even though the effects of context on certain judgments and behaviors has been extensively shown, it has not yet been determined if these contrast effects can affect *trust* behaviors, particularly in non-human systems (automation).

Automation and Trust. It has been proposed by Parasuraman, Sheridan, and Wickens (2000) that there are four stages of automation that roughly correspond to the four stages of human information processing. The first two stages usually involve stimulus input and diagnosis, whereas the latter two involve response selection and execution. Each stage of automation can provide benefits to overall human-automation performance, depending on the task at hand.

Although some research has analyzed human-automation interaction in the latter two stages (e.g. Rovira, McGarry & Parasuraman, 2007; Sarter & Schroeder, 2001), most studies have focused on the early two stages (e.g. Fisher & Tan, 1989; Galster, Bolia, Roe, & Parasuraman, R., 2001; Horrey & Wickens, 2001; Merlo, Wickens, & Yeh, 2000; Wickens, Kroft, & Yeh, 2000; Yeh & Wickens, 2001; Yeh, Wickens, & Seagull, 1999), with particular focus on diagnostic automation (e.g. Dixon & Wickens, 2006; Dixon, Wickens & Chang, 2005; Dixon, Wickens & McCarley, 2007; Dzindolet, Pierce, Beck, & Dawe, 1999; Mosier, Skitka, Heers, & Burdick, 1998; Rice & McCarley, 2008; Skitka, Mosier, Burdick, & Rosenblatt, 2000; Wickens & Dixon, 2007). For the current study, we limit our focus to diagnostic automation.

Imperfect Diagnostic Automation. Although diagnostic aids are often beneficial when they are perfectly reliable (Dixon, Wickens & Chang, 2005; see Yeh & Wickens, 2001, for an exception), it is commonly the case that the aid is imperfect, or even unreliable. Imperfect automation occurs when the aid produces errors in diagnostics. For example, a low-fuel light may come on in a vehicle when in fact plenty of fuel remains. Or the traffic collision avoidance system (TCAS) in an aircraft might fail to detect another aircraft encroaching into its airspace.

Viewed from a signal detection theory framework (Green & Swets, 1966), imperfect diagnostic automation errs in one of two ways. It either presents false alarms (FA), whereby the diagnostic aid declares an event when there is none, or it presents a miss, whereby the diagnostic

aid fails to detect an event. In terms of bias, FA-prone automation would be tantamount to a liberal criterion, whereas miss-prone automation would be tantamount to a conservative criterion.

The penalty for each type of diagnostic error depends on the situation at hand and the cost of the failure (Wickens & Xu, 2002). In one tragic example, an automation miss was partially responsible for a Northwest Airlines MD-80 crash on takeoff in Detroit. The automation failed to notify the pilot of incorrectly configured flaps and slats (National Transportation Safety Board, 1988). Another example could be when a target-cuing aid in an unmanned aerial vehicle (UAV) fails to highlight an enemy target during combat. These types of errors can lead to great loss of life, and designers often set the automation criterion in these systems to reflect a liberal bias (Parasuraman, Hancock, & Olofinboba, 1997).

In contrast, FAs may sometimes be as costly as misses. Chiechio, Pritchett, van Paassen, Mulder, Roeden, & Kalaver (2003) discovered that in-flight FAs confused pilots to the degree that the majority of pilots felt something was amiss, but did not know exactly what was wrong. Their response was to make unnecessary maneuvers such as aborted landings. Bliss (2003) argued that FAs often contributed to more aviation problems than did misses. Several experimental studies have found that automation FAs are more deleterious to overall human-automation performance than are automation misses (Cotté, Meyer & Coughlin, 2001; Dingus, McGehee, Manakkal, Jahns, Carney, & Hankey, 1997; Dixon & Wickens, 2006; Dixon, Wickens & McCarley, 2007; Gupta, Bizantz & Singh, 2001; Lehto, et al., 2000; Meyer and Ballas, 1997).

Meyer (2001; 2004) proposed that automation FAs and misses have fundamentally different effects on operator trust. He differentiated between operator compliance and reliance, which are affected by FAs and misses, respectively. Compliance refers to operator trust when the

automation provides an alert. Compliance decreases as the rate of FAs increases. Once the FA-rate reaches a certain threshold, the cry-wolf effect occurs (Breznitz, 1983). Reliance refers to operator trust when the automation is silent, or indicates a non-alert. As the miss-rate increases, reliance decreases. Although some studies have shown that these two constructs are not entirely independent of each other (Dixon & Wickens, 2006; Dixon, Wickens & McCarley, 2007; Wickens, Dixon, Goh & Hammer, 2005), the idea that each type of error has stronger selective effects than non-selective effects appears to be well-founded; that is, an increase in FA rate tends to more negatively affect operator compliance than reliance, whereas an increase in miss rate tends to more negatively affect operator reliance than compliance (Rice & McCarley, in press).

When diagnostic aids produce FAs or misses, the degradation in trust can have negative effects on operator dependence. Trust and dependence are distinct from each other. Trust is a subjective psychological state, whereby the operator puts a certain amount of faith in a particular system. Dependence is the objective behavioral result, which can be measured by assessing agreement rates, response times, etc. (e.g. Dixon & Wickens, 2006). Trust and dependence do not necessarily correlate; an example being when a person depends on a piece of computer software to function, but does not particularly trust that program. In this case, necessity drives dependence in the absence of trust. Similarly, one can trust an aid, while not depending on it, particularly if the person does not need the reduction in workload or feels they can outperform the automation. Trust and dependence have varying degrees of calibration with the true reliability of the diagnostic aid, and can lead to various states of overtrust, undertrust, or calibrated trust (Parasuraman & Riley, 1997).

The true reliability of the diagnostic aid refers to the actual performance of the aid, which can be characterized by declaring a probability between 0 and 1 [i.e. $R = 1 - p(\text{Failure})$].

Perfectly reliable automation is given a value of $R = 1.0$. Anything less is imperfect. Clearly, operator trust is impacted by the reliability level, which in turn, can affect operator dependence. If operator trust falls too low, then the diagnostic aid may be abandoned altogether while the operator “goes baseline;” that is, the operator completes the task and totally ignores the aid’s recommendations. There is no absolute line of reliability below which the operator should go baseline; however, some have described a 70% threshold, in which a diagnostic aid that is less reliable than 70% tends to result in human-automation performance that is worse than baseline (Wickens & Dixon, 2007).

Current Study. As reviewed earlier, affect is capable of influencing behavior. However, whether or not it is capable of influencing human *trust* behavior with automation is unclear. To investigate this possibility, participants were first presented with either 9 positive images of automation (e.g. expensive car), 9 negative images of automation (e.g. beat-up car), or no images at all. To insure that the pictures used would truly cause positive or negative affect, a pilot study was conducted. The pilot study presented participants with 18 pictures of automation and were asked to rate each picture on how the automation made them feel. This demonstrated that 9 pictures were given positive affect ratings and the other 9 pictures were given negative affect ratings ($p < .0001$). Therefore, the 9 pictures that were given the positive ratings were used for the Positive automation experimental condition and the 9 pictures that were given the negative ratings were used for the Negative automation experimental condition.

Following the presentation of the Positive or Negative automation pictures or no pictures, participants conducted a visual search task similar to what UAV operators do during combat identification tasks. This task involved searching for enemy tanks in aerial photographs of Baghdad. The task was augmented by the presence of diagnostic automation that gave an

accurate recommendation to participants on each trial as to whether or not it detected an enemy tank. Trust behaviors were measured by analyzing the agreement rates between participants' choices and automation recommendations (higher agreement indicates higher dependence) and response times (RTs) for each trial (quicker RTs indicate high dependence). Presumably, if one has total trust in the automation, then one should always agree quickly with the automation.

METHOD

Participants. Fifty-four NMSU students (35 females) took part for course credit. All participants had normal or corrected to normal vision.

Apparatus and Stimuli. A Dell computer with a 20" monitor was used to display the experiment, which was programmed via E-Prime 1.1. Eighteen (9 positive and 9 negative) images of automation were pre-selected and independently rated ($p < .0001$). One hundred aerial images of Baghdad were created using GoogleEarth. Fifty of the images included a superimposed tank, while the other 50 images contained no tank. The tank took up approximately 2 degrees of visual angle.

Procedure. Participants were seated 21" from the monitor, with distance controlled by a chin rest. The 9 positive or negative images were first presented for 5 seconds each. In the control condition, participants did not view pictures. Following the presentation of the images, participants were given instructions on-screen regarding the experimental task. They were instructed to search each image for an enemy tank and push the J key if they believed a tank was present and the F key if they believed that a tank was not present in the. Participants were told that the images would remain on the screen until they responded. Following each response, feedback was provided as to whether they were correct or not.

Participants were also told that a diagnostic aid would provide recommendations during each trial. The participants were instructed that they could agree or disagree with the automation however they wished. They were not told how reliable the automation was (it was 100% reliable).

Each trial began with the automation recommendation being presented for 1000 ms, followed by the aerial image display. This display remained until the participant responded. Following this, a feedback display was presented for 1000 ms. After 100 trials, the experiment ended automatically.

RESULTS

Response Times (RT). Data for RTs were collapsed into 10 epochs (10 trials per epoch) for parametric analyses (Chun & Jiang, 1998). These data are presented in Figure 1.

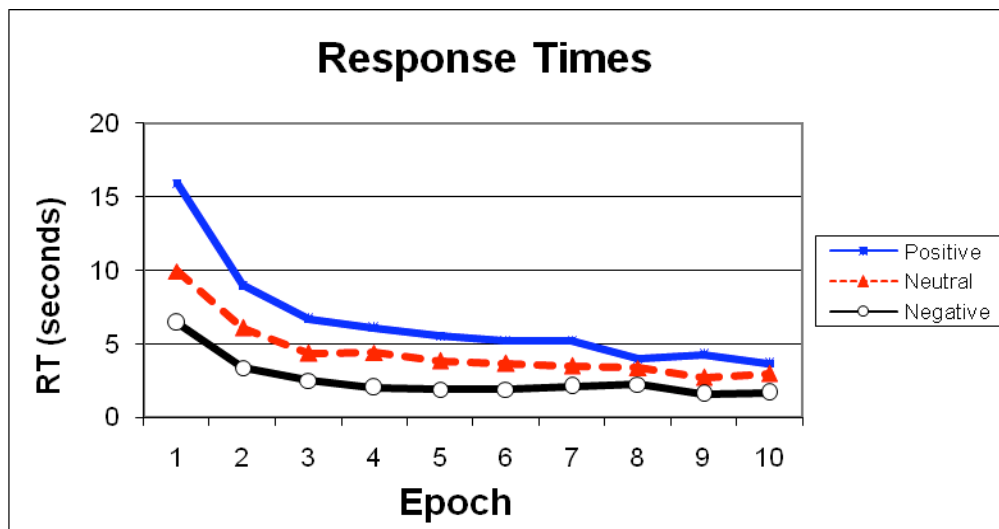


Figure 1. Response times by epoch.

Parametric Analyses. A two-way ANOVA with Condition and Epoch as factors revealed a main effect of Condition, $F(2, 51) = 6.27, p < .01, d = .70$, and a main effect of Epoch, $F(9, 459) = 26.44, p < .001, d = .48$. An interaction between Condition and Epoch, $F(18, 459) = 2.03,$

$p < .01$, $d = .13$, indicated that the differences between conditions dissipated somewhat, but did not disappear, over the course of the experiment. Post hoc tests revealed significant differences between the experimental conditions for all epochs ($ps < .05$). The Neutral condition fell squarely between the other two conditions.

Non-Parametric Analyses. To ensure that we did not capitalize on a few extreme scores that might have accentuated the differences in means, the following binomial comparisons were conducted: a) Neutral vs. Positive, b) Neutral vs. Negative, and c) Positive vs. Negative. The Positive condition resulted in longer RTs relative to the Neutral condition 85% of the time, $p < .000001$. The Negative condition resulted in quicker RTs relative to the Neutral condition 98% of the time, $p < .000001$. The Positive condition resulted in longer RTs relative to the Negative condition 100% of the time, $p < .000001$.

Accuracy. Accuracy measures determined how often participants were correct during the 100 trials. These data are presented in Figure 2.

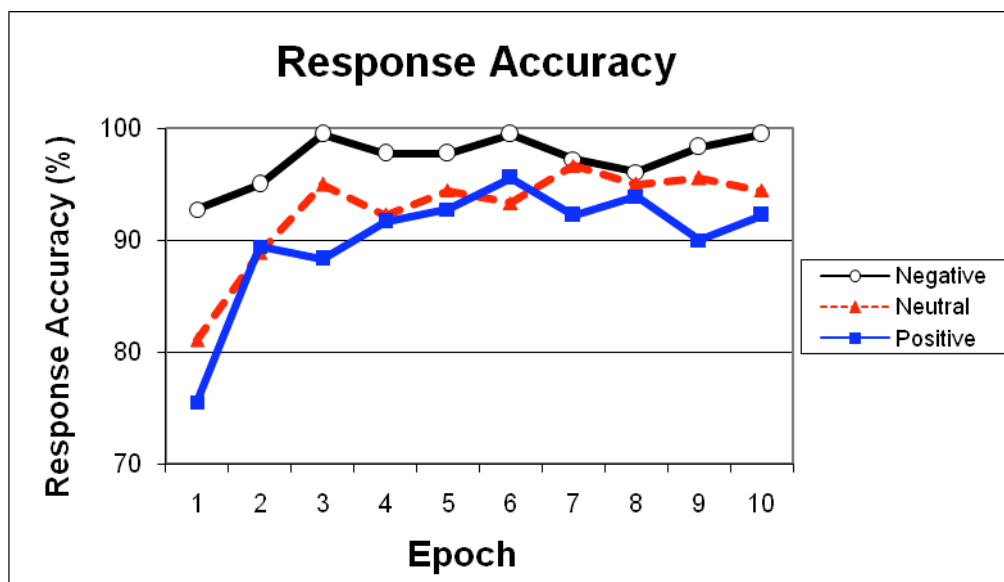


Figure 2. Response accuracy by epoch.

Parametric Analyses. A two-way ANOVA with Condition and Epoch as factors revealed a marginally significant main effect of Condition, $F(2, 51) = 2.52, p = .09, d = .44$, a main effect of Epoch, $F(9, 459) = 9.88, p < .001, d = .29$, and no significant interaction between Condition and Epoch, $F(18, 459) = 1.37, p > .10, d = .11$. While not statistically significant, Figure 2 clearly shows that the differences between the experimental conditions were in the predicted contrasting direction. A post hoc comparison between these two conditions confirmed this effect, $t(34) = 2.16, p < .02, d = .74$.

Non-Parametric Analyses. As with the RT data, the following binomial comparisons were conducted: a) Neutral vs. Positive, b) Neutral vs. Negative, and c) Positive vs. Negative. The Positive condition resulted in lower accuracy relative to the Neutral condition 78% of the time, $p < .000001$. The Negative condition resulted in marginally higher accuracy relative to the Neutral condition 68% of the time, $p = .06$. The Positive condition resulted in lower accuracy relative to the Negative condition 94% of the time, $p < .000001$.

DISCUSSION

The data confirm that not only are trust attitudes easily manipulated by prior priming with positive or negative exemplars, but also that dependence on an automated aid can be manipulated. Both RTs and accuracy were strongly affected by the primes, with clear contrasting effects seen for both experimental conditions. Specifically, prior positive primes resulted in longer RTs and lower overall accuracy, whereas prior negative primes resulted in quicker RTs and higher overall accuracy.

These data demonstrate that priming operators with positive or negative images can have strong effects on their trust in and dependence on different types of automation. Counter-

intuitively, the primes created contrasting effects in trust and dependence, with positive primes resulting in reduced trust and dependence, and negative primes resulting in increased trust and dependence. While it is not overly surprising to find contrasting effects in attitudes (Bodenhausen, Schwarz, Bless, & Wanke, 1995; Bohner, Ruder, & Erb, 2002; Pesta, Dunegan, & Hrivnak, 2007; Tormala & Clarkson, 2007), it is a novel finding to see contrasting effects in repeated behaviors, especially in an extended and complex combat identification task such as the one used in the current study. From a practical standpoint, contrast is clearly not trivial and the effects seen in the current study were not minor. RTs differed by 5-10 seconds and accuracy differed by as much as 10-15%. Effects of this size, if seen in the real world, would have important ramifications in operator training.

The biggest question perhaps is whether to prime operators with positive or negative images. The answer to this question is that it depends on whether a designer wants operators to have higher or lower trust in, and dependence on, the automation. This, of course, depends on how reliable the automation is. If the automation is highly reliable, and particularly if it is much more reliable than the unaided human, then it makes sense to strive for higher dependence. If the automation is more unreliable, then the opposite approach is more practical. Once the goals of the designer and the reliability of the automation are determined, the proper prime could be implemented.

One way to accomplish priming without undo interference is through presenting operators with subliminal primes. By doing this, operators would not be distracted with prime presentations, and if the primes were presented in intervals throughout the task, the effects would be less likely to decay. For instance, UAV operators often have low dependence on automation mainly for the reason that alarms usually alert the controller to a problem that they have already

planned ahead for. Because of this lack of dependence, operators sometimes miss a detrimental warning which in some cases, can lead to a more serious problem than what originally triggered the alert. By frequently priming operators with subliminal pictures of poor automation, it is possible that they will be more trusting of the alarms and therefore, increase overall human-automation performance. Even though this paradigm seems reasonable on the surface, there are caveats. It might be that eventually, operators will begin to recognize or actually see the subliminal primes. If this were to occur, operators might become distracted or curious as to why pictures are frequently appearing on their screen. Much care should be taken to avoid this type of task interference.

It is also possible that UAV operators are already unexpectedly primed with poor or good automation before they begin their task. We are constantly interacting with automation, and sometimes, it behaves poorly. If an operator uses a computer to check their email during a lunch break, and the computer crashes, that operator could now have a negative idea of automation in his head. This negative idea could lead to more trust in the automation involved in his job. This increased trust may or may not be beneficial, depending on the reliability of the aid and the task at hand. On the downside, if an operator drives a very nice vehicle to work every day, her idea of automation could be positive which, according to our data, could lead to a decrease in dependence. Again, depending on the aid and the task, this could be advantageous or disadvantageous. Although these singular instances of poor or good automation interaction might not be enough to influence the operator for the majority of his or her task, it could be enough to influence the first few moments of the task, which in the real world, could have detrimental effects.

Thus, future research should consider both the feasibility of employing priming mechanisms as a way to manipulate operator trust, as well as the unplanned priming effects that occur during daily life. We believe both of these issues warrant the need for further priming and automation research.

Acknowledgments

The authors wish to thank Amy Wells, Gayle Hunt, and Jackie Chavez for their help in collecting data. 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.

References

- Aarts, H. & Dijksterhuis, A. (2002). Category activation effects in judgment and behavior: The moderating role of perceived comparability. *British Journal of Social Psychology, 41*, 123-138.
- Becker, G. A. & Miller, C. E. (2002). Examining contrast effects in performance appraisals: Using appropriate controls and assessing accuracy. *Journal of Psychology: Interdisciplinary & Applied, 136*, 667-683.
- Bliss, J. (2003). An investigation of alarm related accidents and incidents in aviation. *International Journal of Aviation Psychology, 13*, 249-268.
- Bodenhausen, G. V., Schwarz, N., Bless, H., & Wanke, M. (1995). Effects of atypical exemplars on racial beliefs: Enlightened racism or generalized appraisals? *Journal of Experimental Social Psychology, 31*, 48-63.
- Bohner, G., Ruder, M., & Erb, H. P. (2002). When expertise backfires: Contrast and assimilation effects in persuasion. *British Journal of Social Psychology, 41*, 495-519.
- Boninger, D. S., Gleicher, F., & Strathman, A. (1994). Counterfactual thinking: From what might have been to what may be. *Journal of Personality and Social Psychology, 67*, 297-307.
- Brewer, N. T. & Chapman, G. B. (2003). Contrast effects in judgments of health hazards. *Journal of Social Psychology, 143*, 341-354.
- Breanitz, S. (1984). *Cry Wolf: The Psychology of False Alarms*. Englewood Hills, N.J.: Lawrence Erlbaum Associates.
- Brown, D. R. (1953). Stimulus-similarity and the anchoring of subjective scales. *American Journal of Psychology, 66*, 199-214.

- Chieccchio, J., Pritchett, A., van Paassen, M.M., Mulder, M., Roeden, M., & Kalaver, S. (2003). Aiding pilots in detecting faults in the flight control loop. *Proceedings of the 12th Annual International Symposium on Aviation Psychology* (pp. 243-246). Columbus, OH: The Ohio State University.
- Chun, M. M., & Jiang, Y. (1998). Contextual cuing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36, 28-71.
- Cotté, N., Meyer, J., & Coughlin, J. F. (2001). Older and younger driver's reliance on collision warning systems. *Proceedings of the 45th Annual Meeting of the Human Factor Society* (pp. 277-280). Santa Monica, CA: Human Factors and Ergonomics Society.
- Dijksterhuis, A., Spears, R., & Lépinasse, V. (2001). Reflecting and deflecting stereotypes: Assimilation and contrast in impression formation and automatic behavior. *Journal of Experimental Social Psychology*, 37, 286-299.
- Dijksterhuis, A., Spears, R., Postmes, T., Stapel, D. A., Koomen, W., van Knippenberg, A., & Scheepers, D. (1998). Seeing one thing and doing another: Contrast effects in automatic behavior. *Journal of Personality and Social Psychology*, 75, 862-871.
- Dingus, T. A., McGehee, D. V., Manakkal, N., Jahns, S. K., Carney, C., & Hankey, J. M. (1997). Human factors field evaluation of automotive headway maintenance/collision warning devices. *Human Factors*, 39, 216-229.
- Dixon, S. R. & Wickens, C. D. (2006). Automation reliability in unmanned aerial vehicle control: A reliance-compliance model of automation dependence in high workload. *Human Factors*, 48, 474-486.
- Dixon, S.R., Wickens, C.D., & Chang, D. (2003). Comparing quantitative model predictions to experimental data in multiple-UAV flight control. *Proceedings of the 47th Annual Meeting*

- of the Human Factors and Ergonomics Society. Santa Monica, CA: Human Factors and Ergonomics Society.*
- Dixon, S. R., Wickens, C. D., & McCarley, J. S. (2007). On the independence of compliance and reliance: Are automation false alarms worse than misses? *Human Factors, 49*, 564-572.
- Dzindolet, M. T., Pierce, L. G., Beck, H. P., & Dawe, L. A. (1999). Misuse and disuse of automated aids. *Proceedings of the 43rd Annual Meeting of the Human Factors and Ergonomics Society (pp. 339-343). Santa Monica, CA: Human Factors and Ergonomics Society.*
- Fisher, D. L. & Tan, K. C. (1989). Visual displays: Highlighting paradox. *Human Factors, 31*, 17-30.
- Galster, S. M., Bolia, R. S., Roe, M. M., & Parasuraman, R. (2001). Effects of automated cueing on decision implementation in a visual search task. *Proceedings of the 45th Annual Meeting of the Human Factor Society (pp. 321-325). Santa Monica, CA: Human Factors and Ergonomics Society.*
- Geers, A. L. & Lassiter, G. D. (2005). Affective assimilation and contrast: Effects of expectations and prior stimulus exposure. *Basic & Applied Social Psychology, 27*, 143-154.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: John Wiley and Sons.
- Gupta, N., Bisantz, A. M., & Singh, T. (2001). Investigation of factors affecting driver performance using adverse condition warning systems. *Proceedings of the 45th Annual Meeting of the Human Factor Society (pp. 1699-1703). Santa Monica, CA: Human Factors and Ergonomics Society.*

- Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology, 19*, 323-340.
- Hicklin, S. K. & Wedell, D. H. (2007). Learning group differences: Implications for contrast and assimilation in stereotyping. *Social Cognition, 25*, 410-454.
- Horrey, W. J. & Wickens, C. D. (2001). Supporting situation assessment through attention guidance: A cost-benefit and depth of processing analysis. *Proceedings of the 45th Annual Meeting of the Human Factors and Ergonomics Society. Santa Monica, CA: Human Factors and Ergonomics Society.*
- Jones, A. M. & Buckingham, J. T. (2005). Self-esteem as a moderator of the effect of social comparison on women's body image. *Journal of Social and Clinical Psychology, 24*, 1164-1187.
- Landman, J., Vandewater, A. A., Stewart, A. J., & Malley, J. E. (1995). Missed opportunities: Psychological ramifications of counterfactual thought in midlife women. *Journal of Adult Development, 2*, 87-97.
- Lehto, M.R., Papastavrou, J.D., Ranney, T.A., & Simmons L.A. (2000). An experimental comparison of conservative versus optimal collision avoidance warning system thresholds. *Safety Science, 36*, 185-209.
- Merlo, J. L., Wickens, C. D., & Yeh, M. (2000). Effect of reliability on cue effectiveness and display signaling. *Proceedings of the 4th Annual Army Federated Laboratory Symposium (pp. 27-31). College Park, MD: Army Research Federated Laboratory Consortium.*
- 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
- Meyer, J., & Ballas, E. (1997). A two-detector signal detection analysis of learning to use alarms. *Proceedings of the 41st Annual Meeting of the Human Factor Society* (pp. 186-189). Santa Monica, CA: Human Factors and Ergonomics Society.
- Mosier, K. L., Skitka, L. J., Heers, S., & Burdick, M. (1998). Automation bias: Decision making and performance in high-tech cockpits. *The International Journal of Aviation Psychology*, 8, 47-63.
- National Transportation Safety Board. 25 March 2008.
<http://www.nts.gov/nts/AccList.asp?month=11&year=1988>.
- Parasuraman, R., Hancock, P.A., & Olofinboba, O. (1997). Alarm effectiveness in driver-centered collision-warning systems. *Ergonomics*, 40, 390-399.
- Parasuraman, R. & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39, 230-253.
- Parasuraman, Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics. Part A: Systems and Humans*, 30, 286-297.
- Pesta, B. J., Dunegan, K. J., & Hrivnak, M. W. (2007). Contrast and rater-perspective effects on judgments of sexual harassment severity: What he thinks she thinks, and vice versa. *Journal of Business and Psychology*, 22, 155-165.
- Rice, S. & McCarley, J. (2008). The Effects of Automation Bias and Saliency on Operator Trust. *International Congress of Psychology*.
- Roese, N.J. (1997). Counterfactual thinking. *Psychological Bulletin*, 121, 133-148.
- Rovira, E., McGarry, K., & Parasuraman, R. (2007). Effects of imperfect automation on decision making in a simulated command and control task. *Human Factors*, 49,76-87.

- Sarter, N. B. & Schroeder, B. (2001). Supporting decision making and action selection under time pressure and uncertainty. The case of in-flight icing. *Human Factors*, *43*, 573-585.
- Schubert, T. W. & Häfner, M. (2003). Contrast from social stereotypes in automatic behavior. *Journal of Experimental Social Psychology*, *39*, 577-584.
- Skitka, L. J., Mosier, K. L., Burdick, M., & Rosenblatt, B. (2000). Automation bias and errors: Are crews better than individuals? *The International Journal of Aviation Psychology*, *10*, 85-97.
- Stapel, D. A. & Koomen, W. (1997). Using primed exemplars during impression formation: Interpretation or comparison? *European Journal of Social Psychology*, *27*, 357-367.
- Stapel, D. A. & Winkielman, P. (1998). Assimilation and contrast as a function of context-target similarity, distinctness, and dimensional relevance. *Personality & Social Psychology Bulletin*, *24*, 634-647.
- Sumer, H. C. & Knight, P. A. (1996). Assimilation and contrast effects in performance ratings: Effects of rating the previous performance on rating subsequent performance. *Journal of Applied Psychology*, *81*, 436-442.
- Tennis, G. H. & Dabbs, J. M. (1975). Judging physical attractiveness: Effects of judges' own attractiveness. *Personality and Social Psychology Bulletin*, *1*, 513-516.
- Tormala, Z. L. & Clarkson, J. J. (2007). Assimilation and contrast in persuasion: The effects of source credibility in multiple message situations. *Personality and Social Psychology Bulletin*, *33*, 559-571.
- Walchli, S. B. & Landman, J. (2003). Effects of counterfactual thought on postpurchase consumer affect. *Psychology & Marketing*, *20*, 23-46.
- Wickens, D. C. & Dixon, S. (2007). The benefits of imperfect diagnostic automation: A synthesis of the literature. *Theoretical Issues in Ergonomic Science*, *8*, 201-212.

- Wickens, C.D., Dixon, S.R., Goh, J. & Hammer, B. (2005). Pilot dependence on imperfect diagnostic automation in simulated UAV flights: an attentional visual scanning analysis. *In Proceedings of the 13th Annual International Symposium of Aviation Psychology*. Dayton, Ohio.
- Wickens, C. D., Kroft, P., & Yeh, M. (2000). Data base overlay in electronic map design: Testing a computational model. *Proceedings of the IEA 2000/HFES 2000 Congress*. Santa Monica, CA: Human Factors & Ergonomics Society.
- Wickens, C. D., & Xu, X. (2002). Automation trust, reliability and attention HMI 02-03 (AHFD-02-14/MAAD-02-2). Savoy, IL: University of Illinois.
- Yeh, M., & Wickens, C. D. (2001). Display signaling in augmented reality: The effects of cue reliability and image realism on attention allocation and trust calibration. *Human Factors*, 43, 355-365.
- Yeh, M., Wickens, C. D., & Seagull, F. J. (1999). Target cueing in visual search: The effects of conformality and display location on the allocation of visual attention. *Human Factors*, 41, 524-542.