

Examining Trust and Reliance in Collaborations between Humans and Automated Agents

J. S. Elson
University of Nebraska
at Omaha
jselson@unomaha.edu

Douglas C. Derrick
University of Nebraska
at Omaha
dcderrick@unomaha.edu

Gina S. Ligon
University of Nebraska
at Omaha
gligon@unomaha.edu

Abstract

Human trust and reliance in artificial agents is critical to effective collaboration in mixed human computer teams. Understanding the conditions under which humans trust and rely upon automated agent recommendations is important as trust is one of the mechanisms that allow people to interact effectively with a variety of teammates. We conducted exploratory research to investigate how personality characteristics and uncertainty conditions affect human-machine interactions. Participants were asked to determine if two images depicted the same or different people, while simultaneously considering the recommendation of an automated agent. Results of this effort demonstrated a correlation between judgements of agent expertise and user trust. In addition, we found that in conditions of high and low uncertainty, the decision outcomes of participants moved significantly in the direction of the agent's recommendation. Differences in reported trust in the agent were observed in individuals with low and high levels of extraversion.

1. Introduction

Many people interact with automated agents every day (e.g., Apple's Siri, Amazon's Alexa, virtual customer-service agents, etc.), and decision-makers at all levels of organizations utilize automated systems that are designed to enable better, faster, and more effective decisions. Human-machine teams are increasingly common as the volume of data and data sources make it impossible for a human to capture and process all available and relevant data. Our reliance on and trust in automated technologies is changing the way we process information [12], decide [18], and act [20]. Understanding the conditions under which humans trust and rely upon automated agents is important, as trust is one of the mechanisms that allow

humans to interact effectively with a variety of teammates [30, 49].

While, there is a large body of research in the social sciences regarding the nature of interpersonal trust, it is one thing to say that we trust a person and something rather different to say that we trust a machine. Machines range in "ability" from calculators to sophisticated artificial intelligence, but ultimately are all non-feeling, non-living partners we rely upon to make decisions. Advancements in information technology, including artificial intelligence, allows humans to work alongside computers as teammates [49]. In a meta-analysis of the effects of interpersonal trust on task performance, Colquitt, Scott, and LePine [14] found that trust in teammates predicted important organizational outcomes such as affective commitment, citizenship behaviors, and task performance in general. This may be because once an individual trusts a teammate, he or she is no longer bound by observing interactions to try to understand motivation and accuracy, which frees up cognitive space and working memory to focus on the task at hand. McEvily, Perrone, and Zaheer [35] found that trust is the underlying mechanism that increases overall effort that individuals apply to group tasks, as well as the extent to which they collaborate toward shared goals.

The need to explore the conditions surrounding trust in human-machine teams was observed in our partnership to design and deploy an automated system used by decision makers in a Fortune 500 company. Financial agreement professionals utilize the information system to make more than 40,000 decisions each month relating to a variety of revenue collection processes. These individuals vary on a number of individual difference including extraversion, a personality trait shown to be important in human-computer teams [7, 48]. Understanding the factors surrounding how individual decision makers judge the expertise of their computer teammates as well as act upon information they are provided is critical to

informing both system design and customized training plans for new system orientation. Where individual personality traits are difficult to change and stable over time [44], tailoring agent attributes to individual users or classes/types of users may provide ways to improve decision making and enhance cooperation between humans and computers.

In this present effort, we conducted exploratory research to examine the individual differences and conditions relating to an individual's trust or distrust of the machine's recommendations. Intelligent decision aides are becoming a more common component of information systems, used across industries and domains to assist decision makers. Understanding the way individuals collaborate with a decision aid will greatly improve how system designers and human computer interaction scientist understand the implication for integrating automated partners into existing teams and workflows. It has been shown that individuals will often defer to a decision aid, but the conditions and circumstances surrounding this phenomenon are not well understood. In this endeavor, we reviewed relevant literature and conducted an exploratory study to determine how a key personality characteristic and uncertainty affect human-machine interactions.

2. Background

There are three primary theoretical frameworks that are utilized in this study: intelligent decision aids, trust and expertise, and individual characteristics. The decision aids and the personality of the individual collaborating with that aid interact to impact trust, reliance, and outcomes. We will briefly review each of these frameworks and tie them together in our experiment.

2.1. Intelligent Decision Aids and Bias

An intelligent decision aid, also referred to as an expert system in some literature, utilizes expert knowledge encapsulated into a software system to provide meaningful information output that can be leveraged by decision makers to assist in cognitive task completion [4]. In this work, we view intelligent decision aids as a technology laying at the cross section of automated agents and decision support systems (see Figure 1). In the book *The Media Equation*, Reeves and Nass [43] provide evidence that human's may perceive computers as social actors and often treat them like they would treat living people. We propose that a similar phenomenon occurs with human users of intelligent decision aids. The study of intelligent

decision aids and the perception of these as intelligent or automated agents is important as these are used in wide variety of areas including systems that aid in making business or financial analysis decisions [4], medical diagnostic and patient care recommendation systems [1], and military strategic decision support systems [42].

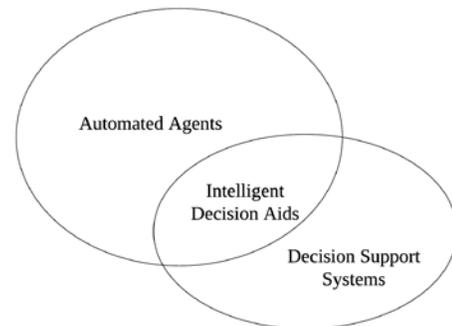


Figure 1. Intelligent Decision Aids

Individuals may encounter intelligent decision aids through interaction with numerous software interfaces or specially designed information dashboards used to highlight strategic information deemed most useful for a user environment. The output of intelligent decision aids can vary tremendously yet are usually comprised of variations around three key characteristics including: form (the numerical, verbal, or pictorial representation of information), organization (grouping, hierarchy, and pattern emphasis of information), and sequence (order and arrangement of information) [28]. The decision of how to present information from an intelligent decision aid is often matched with either the specific task or operational environment in which the system is being used.

Decision aids, such as information displays have been shown to assist decision makers in making better decisions [28]. It has also been observed that intelligent decision aids may outperform humans on a number of judgment tasks [27]. Because intelligent decision aids are not prone to the shortcomings of human cognitive processing, it is often advantageous to rely upon the recommendation or analysis of a well-designed intelligent system. However, human decision making can be biased when supported by imperfect automated decision aids. This bias can be benign or potentially beneficial when a decision aid provides correct recommendations (e.g., by speeding up decision making). The tendency of a human to defer to a decision aid however is more problematic if the decision aid is wrong, resulting in omission and commission errors. Several theories exist as to why individuals bias toward recommendations of intelligent

decision aids. Early work examining the use of intelligent decision aids found that the level of expertise of human operators may also be an important factor when considering an individual's propensity to defer judgment to a decision aid [3]. Individuals who are experts or perceive themselves to have expertise utilize decision aides to a lesser extent than non-experts.

Additional research has supported findings that individuals tend to have a bias toward an automated decision aid even under conditions where the aid was incorrect. This phenomenon was shown to occur in a variety of settings, expertise levels, regardless of direction or training, and shown to impact individuals as well as teams [40]. Deference to a decision aid seemed to depend upon the overall reliability of the aid, the level of automation, and user perception of personal accountability for an outcome. The tendency of individuals to defer to the judgments of a decision aid may relate to processes surrounding psychological and social offloading of task burden onto the machine, thereby reducing the complexity of a decision making task [8, 32]. Thinking requires effort and humans are particularly adept at developing or finding shortcuts to accomplish work in more timely and efficient manners.

2.2. Trust and Expertise

Trust in automated agents and intelligent decision aids is critical as without it, decision makers cannot offload cognitive tasks and must instead exert greater cognitive effort every time a decision must be made. A number of factors need to be considered when investigating human trust in a decision aid including an individual's perception of the aids ability, benevolence, and integrity [33]. In humans, there is evidence suggesting people tend to trust a new relationship from the start as not doing so requires extra cognitive effort [32]. In supervisory control environments however, human operators may have bias toward mistrust of a machine when first beginning interaction with the system [47]. Muir [38] suggested that this contradiction is explained by different levels of risk associated with different human machine interactions. The greater the perceived risk to users of a system, the greater the initial level of mistrust.

Muir [38] looked specifically at trust in decision aids, and suggested that in order to build trust, humans needed to interact with or utilize the system. If an individual distrusted a decision aid, the absence of attention paid to the aid would make positive changes in trust perceptions difficult. The level of confidence an individual places in a decision aid is directly tied to whether or not an individual chooses to use the aid at

all [5]. Individuals who trust a decision aid are thus more likely to use the aid when deciding compared to those individuals who do not. Reliance on a decision aid has been shown to increase when perceptions of initial face validity are high [5]. Therefore, individuals who initially consider a decision aid to be competent or expert would utilize and trust a decision aid more than those who do not. The relationship between trust and perceptions of system expertise should hold up regardless of the form of a decision aid. Research with embodied agents showed user trust were similarly related to perceived expertise of the system [17].

Just like human-to-human trust, experiences with an agent affect perceptions and subsequent interactions. Perceptions of trust are not static and are changing depending upon outcomes, risk and context. Interpersonal Adaptation Theory (IAT) [11] proposes that dyadic communication is adaptive and purposeful and that when engaging in any type of interaction, people enter with their own set of expectations, beliefs, motivations, requirements, and desires. IAT also predicts that time is critical to predicting behavior. Because we are always adapting, our behaviors will be dynamically changing in response to our affect, speaking partner, and environment. It follows that time will naturally affect the nature of trust. Elkins and Derrick [20] showed that for one particular interaction between humans and automated agents, trust was temporally variant and could be predicted by a linear change in time.

2.3. Individual Characteristics

The Five Factor Model or "Big Five" offers a practical arrangement of personality components that have been shown to affect human team performance [46], and these five traits are the primary way to measure job-related personality in the organizational psychology and team literature [6]. In addition, the Big Five have been examined in relationship to job satisfaction, relationship satisfaction, and life satisfaction [2, 9, 13]. The five personality traits include: openness, conscientiousness, extraversion, agreeableness, and neuroticism and have been shown to be very stable over extended periods of time [22]. In this work, we focus on extraversion as it is one of the factors most related to interactions with a variety of teammates [41].

Extraversion is a broad factor that encompasses the tendency to be energetic, affiliative, and dominant [31]. Previous studies have shown a positive correlation between extraversion and the levels of participation in computer-mediated teams [7, 48].

Extraversion also is shown to be related to positive affect [16], which is characteristic of individuals who experience high degree of job satisfaction in what has traditionally been a human-to-human-centered workplace [15]. However, when looking at the underlying facets of the introversion-extraversion scale, such as dominance and sociability [19, 34], introverts, or those low in extraversion, prefer medium-to low-energy situations that do not require them to play a dominant role. Introverts like to perceive and take information in, instead of offering information [19, 34]. Additionally, McKenna and Bargh [36] suggest that anonymity associated with the Internet and online tasks coupled with high feelings of control make interacting with computers ideal for introverts.

3. Conceptual Model and Hypothesis

Research has consistently supported the assertion that trust is multidimensional and consists of many interrelated factors [29]. Some of the factors that affect trust include ability (the group of skills, competencies, and characteristics that enable a person to have influence within some specific domain), benevolence (the extent to which a trustee is believed to want to do good to the trustor) and integrity (the trustor's perception that the trustee adheres to a set of principles that the trustor finds acceptable) [18].

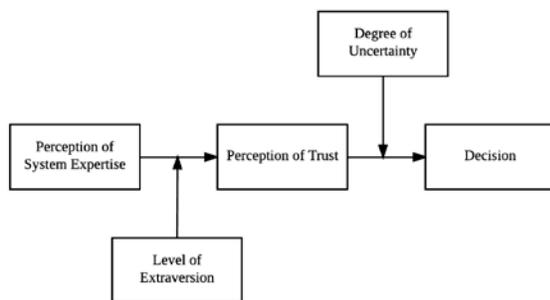


Figure 2. Conceptual Model of Decision Making in Human Computer Teams

Our model is made up of five components (shown in Figure 2). The first was perception of system ability, which consisted of the decision aid indicating correct and incorrect task assessments. Next, we measured perceptions of trust throughout the experiment and participants reported how trustworthy they felt the intelligent decision aid was. Level of extraversion, an individual characteristic of the Five-Factor Model of personality [23], was believed to moderate the relationship between these two variables. The assessment of individual levels of extraversion were collected at a separate time period and completed by all individuals in the participant pool. The decision or

outcome made by each participant was collected throughout the study as individuals complete the assessment task they were assigned. Moderating the relationship between perceptions of trust and decision outcome was a variable relating to the degree of uncertainty about the correct “response” to the task, or decision. Uncertainty was based on responses obtained in a pilot study.

We expect that evaluations of expertise (system ability) will have a direct relationship with users’ trust in their partner. If the aid is seen to be accurate, the perceived trust in the automated partner will increase. It would follow that the inverse also is true, for example if the partner is seen to be inaccurate. In such a case, perceived trust in the partner will decrease. We therefore propose the following hypothesis:

H1: The judgement of expertise of the intelligent decision aid is highly correlated to the perception of trustworthiness of that aid.

Deference to a decision aid by experts can vary depending on task and conditions surrounding the task. In this study, uncertainty was operationalized by the results of a pre-test where another sample of made judgments regarding identity of the same individual represented in two different images. Decision aids may be utilized more when users are facing situations of high uncertainty [10]. In conditions of high uncertainty, we expect humans will defer to the decision of the agent. Therefore, we propose the following hypothesis:

H2: In decision making tasks involving high uncertainty, humans will defer to the automated intelligent decision aid.

When individuals encounter messaging that is contradictory to their own assessments, personality characteristics in combination with the level of uncertainty is believed to influence an individual’s perception of trust. In situations of high uncertainty, it is believed that individuals, regardless of their level of extraversion, will defer to the intelligent decision aid. The relationship between system ability and trust may be moderated by an individual’s personality characteristics. On average, extraverts are more willing to trust other people than are introverts [21]. Given that extraversion is characterized by a high need for affiliation and relationships, individuals who are high in extraversion are also more likely to trust machines [37]. However, Merritt & Ilgen [37] found that those individuals—extraverts with a higher propensity to trust machines—were less likely to use a machine after

they had experience with it performing poorly. In other words, because extraverts may be more likely to attribute affiliative, anthropomorphic characteristics onto the agent when working with it, violations of that trust can be particularly egregious when the agent performs in a substandard way. However, research on extraversion and willingness to trust is scant in human teams as well as in mixed human-agent teams, so we are unable to hypothesize how this individual characteristic will operate. Therefore, we propose the following research question:

RQ: What is the relationship between an individual’s perception of intelligent agent trust and extraversion personality trait?

4. Methodology

An experiment was designed to better understand the interplay between trust, ability, and the impact on an individual’s use of an intelligent decision aid under conditions of uncertainty. We used a facial recognition task as the human visual system is particularly well adapted to identifying visual differences [24], thus this task would be immediately familiar to all participants.

In this study, participants were asked to interact with a mock intelligent decision aid designed to simulate an intelligent decision support system (Refer to Figure 3). Participants used the system to complete a series of tasks where they were asked to determine if two images depicted the same person or different people. In addition to the two images being assessed, participants also see an automated agent’s assessment of the image pairs. Using this information, the participants indicated their decision by selecting a radio button indicating their choice for the two images (i.e., the decision was if the people were identical).



Figure 3. Identification Task

4.1. Study Development

We collected 200 pairs of publicly available images (400 images of 200 people, images for each individual were selected from the public domain). We conducted a pre-test where another sample of participants determined whether an individual and corresponding identification card were the same person for each of the 200 images pairs. In addition to the evaluation task, this set of participants in the pre-test indicated their confidence in their decision for each image pair they assessed. Confidence was measured using the following five item Likert-type scale responses: 1) not-confident, 2) somewhat confident, 3) confident, 4) very confident, 5) extremely confident. Thirty-two people participated in the pilot study and results from the pre-test were used to calculate average confidence ratings for participant’s confidence in their assessment of each image pair. From this confidence rating, we selected a total of 16 image pairs to be used as manipulations in our final experimental design. We selected the eight pairs of images ranked the highest in confidence to represent our low uncertainty condition and eight pairs with the lowest ranked confidence to represent our high uncertainty condition.

The experimental task was developed to expose participants to each of the 200 image pairs (eight rounds of 25 pairs) along with a partner recommendation indicating if the people shown were the same or different (Refer to Figure 4). For this study, the partner recommendation area indicated individuals were “the same” for every image pair except for a total of 8 times. At these times, participants were told that the images represented two different people. These incorrect assessments were shown (after the first 8 images and before the last 8 images) in the second, fourth, fifth and seventh rounds. The incorrect partner assessment consisted of four of the images from the high certainty condition and four images from the low certainty condition.

Round	Intelligent Decision Aid Recommendation	Uncertainty Condition
1	Same	Low
2	Different	High
3	Same	High
4	Different	High
5	Different	Low
6	Same	Low
7	Different	Low
8	Same	High

Figure 4. Study Task Manipulation

4.2 Participants

Students from a Midwestern university were recruited to participate in this study and were compensated with course credit for their involvement. Our final data set eliminated seven individuals (<10% of total sample) who had incomplete answers to our individual characteristics assessment. The final data consisted of 31 males and 33 female subjects. The average age of the male subjects was 23.6, median age was 21 with a standard deviation of 4.7. The average age of the female subjects was 22.6, median age was 21 with a standard deviation of 5.6.

4.3 Procedure

Upon arrival to the lab, each participant was given an initial briefing that gave an overview of the study that was about to be completed. Each participant was given an informed consent packet and given an opportunity review it and ask the researchers any questions they had regarding their participation in the study. After consent was obtained, each participant was seated in front of a dedicated computer terminal that hosted the experiment and positioned in front of a computer monitor with eye-tracking capabilities. At the beginning of each study participants were asked to complete a short calibration process (Tobii regular calibration screen) that required them to focus on 9 dots positioned with three rows of dots across the top, middle, and bottom of the screen. This process was repeated until the participant acquired an Excellent calibration (average distance of measured gaze from the target $\mu(x,y) \leq 20$ pixels) was achieved. Participants were informed to limit unnecessary eye movement and to avoid touching their face with their hands as this would impact the quality of the eye-tracking data. Participants were then read a brief description of the task and were instructed to begin the study.

Participants were exposed to 25 image pairs over the course of 8 rounds. After every round, participants were asked to rate their perception of their partners trustworthiness and expertise using verified semantic differential word pairings [39]. Following the 8 rounds, users were asked to describe their perceptions of their partner (ability, benevolence and integrity) in an open response question format [33]. After completing the study, participants were debriefed and thanked for their participation.

5. Analysis and Results

5.1. Correlation Between Trust and Expertise

We conducted a factor analysis on respondent ratings relating to the system's expertise and trustworthiness. We explored the relationship across the eight rounds participants interacted with the system and found a significant correlation between trust in the system and judgements of expertise. The following shows the results of the correlation between the variables expertise and trustworthiness for each round, N=64, R1=.855, R2=.850, R3=.831, R4=.874, R5=.825, R6=.823, R7=.887, R8=.860.

We then calculated Cronbach's Alpha for each of these rounds and found the following results. For round 1 factor 1: .948 and factor 2: .949. For round 2 factor 1: .962 and factor 2: .917. For round 3 factor 1 : .953 and factor 2: .948. For round 4 factor 1: .955 and factor 2: .953. For round 5 factor 1: .963 and factor 2: .928. For round 6 factor 1: .959 and factor 2: .947. For round 7 factor 1: .969 and factor 2: .957. For round 8 factor 1: .947 and factor 2: .974. These results indicate excellent internal consistency across the semantic differential word pairs used to assess these ratings for all the rounds.

These analysis results of this particular effort suggest that there is a significant correlation between an individual's perception of expertise and trust in the automated decision aid. Under the conditions of this research, individuals who consider their partner to be expert trust the decision aid more than those who do not consider their partner to be an expert. Therefore, Hypothesis 1 received support.

5.2. Decision Outcomes with High Uncertainty

To test the impact of the agent in conditions of highest uncertainty, we analyzed the three stimuli that were rated the most uncertain in the pilot study, and conducted a Chi-square test to compare between the two groups (the first group did not have the help of the agent, N=32) and the second group had the agent recommendation (N=64). We compared the decision (same person, different person) between the two groups for each stimulus and found a significant difference. In each case, the decision made by the participants moved significantly in the direction of the agent's recommendation, regardless if that decision was to recommend that the images were similar or different. In two of the cases, the agent said that the images were the same, and in the other case the agent

said that they were different. The decision outcomes always moved in accordance with the recommendation of the agent. Table 1 shows the results of this analysis where ‘S’ means judgment was the “Same Person” and ‘D’ means judgment was “Different Person”.

Table 1. Decisions With and Without Agent Support in Conditions of High Uncertainty

	Pilot Decisions		With Agent Decisions		Chi-Square	p-value
	S	D	S	D		
Stimulus 1	16	16	46	18	4.46	p < .05
Stimulus 2	16	16	54	10	12.76	p < .001
Stimulus 3	18	14	22	42	4.2	p < .05

The results of this analysis are consistent with literature presented in the background section of this paper. Where decisions in the pilot study were evenly split between the two selection options, the presence of a partner recommendations communicated through a decision aid significantly increased participant responses in both directions in the three high uncertainty conditions. Therefore, hypothesis 2 received support.

5.3. Perceptions of Trust for Individuals of High and Low Extraversion

The individual personality characteristic, extraversion was assessed using the Big Five Index, a 44-item instrument that measures extraversion, agreeableness, openness to experience, conscientiousness, and neuroticism [25, 26]. Participants were instructed to rate the extent to which they agree or disagree with each statement on a scale from one to five, where one is “Strongly Disagree” and five is “Strongly Agree.” Because extraversion was the primary construct interest for the present effort, we examined these scales in most detail for our analyses. The Cronbach’s alphas for extraversion was 0.78, indicating adequate scale reliability. Individual items were averaged together to create a mean score for each participant.

We took individuals who scored in the top and bottom 10% of the extraversion scale and plotted their perceptions of trust as recorded after each of the eight rounds of the image assessment task (See Figure 5).

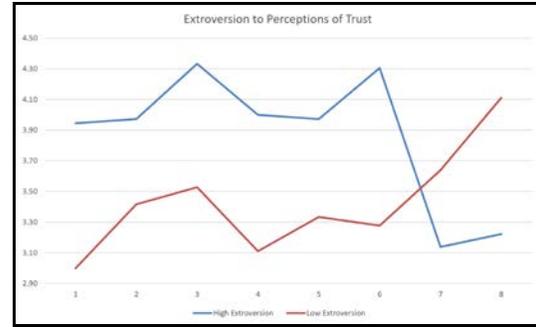


Figure 5. Extraverts and Introverts Trust over Time

Individuals high in extraversion initially rated their trust in the agent as higher than those low in extraversion. This was to be expected as prior studies in human-computer teams have shown extraverts to be more accepting of computer teammates [16]. During the fifth and seventh rounds, the agent gave a recommendation contrary to what was expected in highly confident conditions. After the fifth and seventh rounds, trust in the agent increased compared to the previous round for individual’s low in extraversion while trust in the agent decreased compared to the previous round for individuals high in extraversion. In this research context, extraverts appeared to have their trust most dramatically affected by the conditions where the agent gave an incorrect response in the highly confident condition. This is most clearly seen in the low trust score given immediately after the seventh round from extraverts. Similar to what occurs in social teams, it is possible that extraverts saw this low performance from the intelligent agent as a breach of the psychological contract that exists among team members—an implied obligation of reciprocity and performance [45]. While literature of how trust degrades in response to breaches in psychological contracts is centered on social or human teams, it is reasonable to extend the evaluation of this theory to explain how individual differences such as extraversion are related to automated agent behavior and human-machine trust.

6. Conclusion

The increased use of automated agents is changing the way decisions are made by both individuals and teams. This work attempted to investigate several key areas for human collaboration with intelligent systems, specifically the factors surround human trust relationships with an intelligent decision aid. The results from our study suggest a correlation exists between judgements of expertise (ability) and

perceived trust. Although not a new finding, this is an important contribution to the field as researchers and developers continue to investigate the implications for human-machine interactions. Future work should look more closely at the factors that influence perceptions of expertise, such studies would better allow for automated agents and intelligent decision aids to be designed to intentionally invoke or not invoke perceptions of expertise.

Another contribution of this work was shown in the impact of agent recommendation on human decision outcomes in conditions of high uncertainty. The ability for a decision aid to influence users to change their initial perceptions in tasks like the one performed by individuals in this study illustrate the important role of a machine in a computer human team. Systems should be designed to take into account the tendency of humans to defer to a decision aid in order to mitigate unintentional bias in human decision-maker's analysis of problems and task solutions.

The relationships between the individual personality characteristic of extraversion and perceptions of trust, as observed in this study, helped to illustrate the different ways decision makers trust intelligent systems over time. The reaction by individuals of various levels of extraversion, to system recommendations that may have seemed erroneous, suggest that future systems should account for individual differences when communicating recommendations to teammates. Future work in this area would examine the potential mediating mechanism of breaches of psychological contracts as they relate to resilience to agent errors in mixed teams. Results of this and other efforts will help inform the design of agents that support training and job processes tailored to personality types.

7. References

- [1] AID, P.D. Decision Aids for Patients Considering Options Affecting Cancer Outcomes: Evidence of Efficacy and Policy Implications. (1999).
- [2] Allen, J. and Ferguson, G. Human-Machine Collaborative Planning. *Proceedings of the Third International NASA Workshop on Planning and Scheduling for Space*, (2002), 27–29.
- [3] Arkes, H.R., Dawes, R.M., and Christensen, C. Factors Influencing the Use of a Decision Rule in a Probabilistic Task. *Organizational Behavior and Human Decision Processes* 37, 1 (1986), 93–110.
- [4] Arnold, V., Collier, P.A., Leech, S.A., and Sutton, S.G. Impact of Intelligent Decision Aids on Expert and Novice Decision-Makers' Judgments. *Accounting & Finance* 44, 1 (2004), 1–26.
- [5] Ashton, R.H. Pressure and Performance in Accounting Decision Settings: Paradoxical Effects of Incentives, Feedback, and Justification. *Journal of Accounting Research* 28, (1990), 148.
- [6] Barrick, M.R. and Mount, M.K. The Big Five Personality Dimensions and Job Performance: A Meta-Analysis. *Personnel Psychology* 44, 1 (1991), 1–26.
- [7] Barry, B. and Stewart, G.L. Composition, Process, and Performance in Self-Managed Groups: The Role of Personality. *Journal of Applied Psychology* 82, 1 (1997), 62–78.
- [8] Borum, R. The Science of Interpersonal Trust. (2010).
- [9] Bradshaw, J.M., Feltovich, P.J., Jung, H., Kulkarni, S., Taysom, W., and Uszok, A. Dimensions of Adjustable Autonomy and Mixed-Initiative Interaction. In *Agents and Computational Autonomy*. Springer, Berlin Heidelberg, 2003, 17–39.
- [10] Brown, D.L. and Jones, D.R.. Factors that Influence Reliance on Decision Aids: A Model and an Experiment. *Journal of Information Systems* 12, 2 (1998), 75–94.
- [11] Burgoon, J.K., Stern, L.A., and Dillman, L. *Interpersonal Adaptation: Dyadic Interaction Patterns*. Cambridge University Press, Cambridge, 1995.
- [12] Carr, N. *The Shallows: What the Internet Is Doing to Our Brains*. W. W. Norton & Company, New York, 2011.
- [13] Christoffersen, K. and Woods, D.D. How to Make Automated Systems Team Players. *Advances in Human Performance and Cognitive Engineering Research* 2, (2002), 1–12.
- [14] Colquitt, J.A., Scott, B.A., and LePine, J.A. Trust, Trustworthiness, and Trust Propensity: a Meta-Analytic Test of Their Unique Relationships With Risk Taking and Job Performance. *Journal of Applied Psychology* 92, 4 (2007), 909–927.
- [15] Connolly, J.J. and Viswesvaran, C. The Role of Affectivity in Job Satisfaction: A Meta-Analysis. *Personality and Individual Differences* 29, 2 (2000), 265–281.
- [16] Costa, P.T. and McCrae, R.R. Influence of Extraversion and Neuroticism on Subjective Well-Being: Happy and Unhappy People. *Journal of Personality and Social Psychology* 38, 4 (1980), 668.

- [17] Derrick, D., Jenkins, J., and Jay F. Nunamaker, J. Design Principles for Special Purpose, Embodied, Conversational Intelligence with Environmental Sensors (SPECIES) Agents. *AIS Transactions on Human-Computer Interaction* 3, 2 (2011), 62–81.
- [18] Derrick, D.C. and Ligon, G.S. The Affective Outcomes of Using Influence Tactics in Embodied Conversational Agents. *Computers in Human Behavior* 33, 1 (2014), 39–48.
- [19] Digman, J.M. Personality Structure: Emergence of the Five-Factor Model. *Annual Review of Psychology* 41, 1 (1990), 417–440.
- [20] Elkins, A.C. and Derrick, D.C. The Sound of Trust: Voice as a Measurement of Trust During Interactions with Embodied Conversational Agents. *Group Decision and Negotiation* 22, 5 (2013), 897–913.
- [21] Gaines Jr, S.O., Panter, A.T., Lyde, M.D., et al. Evaluating the Circumplexity of Interpersonal Traits and the Manifestation of Interpersonal Traits in Interpersonal Trust. *Journal of Personality and Social Psychology* 73, 3 (1997), 610.
- [22] Goldberg, L.R. An Alternative Description of Personality: The Big-Five Factor Structure. *Journal of Personality and Social Psychology* 59, 6 (1990), 1216.
- [23] Goldberg, L.R. The Development of Markers for the Big-Five Factor Structure. *Psychological assessment* 4, 1 (1992), 26.
- [24] Hoffman, E.A. and Haxby, J.V. Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nature neuroscience* 3, 1 (2000), 80–84.
- [25] John, O.P., Donahue, E.M., and Kentle, R.L. *The Big Five inventory—Versions 4a and 54*. Berkeley, CA: University of California, Berkeley, Institute of Personality and Social Research, 1991.
- [26] John, O.P., Naumann, L.P., and Soto, C.J. Paradigm Shift to the Integrative Big Five Trait Taxonomy. *Handbook of Personality: Theory and Research* 3, (2008), 114–158.
- [27] Kleinmuntz, B. Why we still use our heads instead of formulas: Toward an integrative approach. *Psychological Bulletin* 107, 3 (1990), 296.
- [28] Kleinmuntz, D.N. and Schkade, D.A. Information Displays and Decision Processes. *Psychological Science* 4, 4 (1993), 221–227.
- [29] Kracher, B., Corritore, C.L., and Wiedenbeck, S. A foundation for understanding online trust in electronic commerce. *Journal of Information, Communication and Ethics in Society* 3, 3 (2005), 131–141.
- [30] Lewis, C.A. Is the Relationship Between Religiosity and Personality “Contaminated” by Social Desirability as Assessed by the Lie Scale? A Methodological Reply to Michael W. Eysenck (1998). *Mental Health, Religion & Culture* 2, 2 (1999), 105–114.
- [31] Lucas, R.E., Diener, E., Grob, A., Suh, E.M., and Shao, L. Cross-Cultural Evidence for the Fundamental Features of Extraversion. *Journal of Personality and Social Psychology* 79, 3 (2000), 452–468.
- [32] Luhmann, N. Trust and power. (1982).
- [33] Mayer, R.C., Davis, J.H., and Schoorman, F.D. An Integrative Model of Organizational Trust. *Academy of Management Review* 20, 3 (1995), 709–734.
- [34] McCrae, R.R. and John, O.P. An Introduction to the Five-Factor Model and its Applications. *Journal of Personality* 60, 2 (1992), 175–215.
- [35] McEvily, B., Perrone, V., and Zaheer, A. Trust as an Organizing Principle. *Organization Science* 14, 1 (2003), 91–103.
- [36] McKenna, K.Y. and Bargh, J.A. Plan 9 from Cyberspace: The Implications of the Internet for Personality and Social Psychology. *Personality and Social Psychology Review* 4, 1 (2000), 57–75.
- [37] Merritt, S.M. and Ilgen, D.R. Not all Trust is Created Equal: Dispositional and History-Based Trust in Human-Automation Interactions. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 50, 2 (2008), 194–210.
- [38] Muir, B.M. Trust Between Humans and Machines, and the Design of Decision Aids. *International Journal of Man-Machine Studies* 27, 5–6 (1987), 527–539.
- [39] Ohanian, R. Construction and Validation of a Scale to Measure Celebrity Endorsers’ Perceived Expertise, Trustworthiness, and Attractiveness. *Journal of advertising* 19, 3 (1990), 39–52.
- [40] Parasuraman, R. and Manzey, D.H. Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 52, 3 (2010), 381–410.
- [41] Peeters, M.A., Rutte, C.G., van Tuijl, H.F., and Reymen, I.M. The Big Five Personality Traits and Individual Satisfaction With the Team. *Small Group Research* 37, 2 (2006), 187–211.

[42] Rasch, R., Kott, A., and Forbus, K.D. Incorporating Ai into Military Decision Making: An Experiment. *IEEE Intelligent Systems* 18, 4 (2003), 18–26.

[43] Reeves, B. and Nass, C. *How People Treat Computers, Television, and New Media Like Real People and Places*. CSLI Publications and Cambridge University Press, New York, 1996.

[44] Roberts, B.W. and DeVecchio, W.F. The Rank-Order Consistency of Personality Traits from Childhood to Old Age: A Quantitative Review of Longitudinal Studies. *Psychological bulletin* 126, 1 (2000), 3.

[45] Rousseau, D.M. Psychological and Implied Contracts in Organizations. *Employee responsibilities and rights journal* 2, 2 (1989), 121–139

[46] Salas, E., Sims, D.E., and Burke, C.S. Is There a “Big Five” in Teamwork? *Small Group Research* 36, 5 (2005), 555–599.

[47] Sheridan, T.B. and Hennessy, R.T. *Research and modeling of supervisory control behavior. Report of a workshop*. DTIC Document, 1984.

[48] Straus, S.G. Getting a Clue The Effects of Communication Media and Information Distribution on Participation and Performance in Computer-Mediated and Face-to-Face Groups. *Small Group Research* 27, 1 (1996), 115–142.

[49] Sycara, K. and Lewis, M. Integrating Intelligent Agents into Human Teams. (2004).