Fiscal Voor:	EV 2022	Task Last Undated	EV 00/06/2022
PI Name	Lee John Ph D	Task Last Opuateu.	1109/00/2022
Project Title	HCAAM VNSCOR. Conversation Analysis to Measure and Manage Trust in Virtual Assistants		
Troject Inter		jois to measure and manage m	
Division Name:	Human Research		
Program/Discipline:			
Program/Discipline Element/Subdiscipline:			
Joint Agency Name:		TechPort:	No
Human Research Program Elements:	(1) HFBP:Human Factors & Behavioral	Performance (IRP Rev H)	
Human Research Program Risks:	(1) HSIA: Risk of Adverse Outcomes Du	e to Inadequate Human System	s Integration Architecture
Space Biology Element:	None		
Space Biology Cross-Element Discipline:	None		
Space Biology Special Category:	None		
PI Email:	jdlee@engr.wisc.edu	Fax:	FY
PI Organization Type:	UNIVERSITY	Phone:	608-890-3168
Organization Name:	University of Wisconsin, Madison		
PI Address 1:	Department of Industrial and Systems Engineering		
PI Address 2:	1513 University Ave		
PI Web Page:			
City:	Madison	State:	WI
Zip Code:	53706-1539	Congressional District:	2
Comments:			
Project Type:	Ground	Solicitation / Funding Source:	2017-2018 HERO 80JSC017N0001-BPBA Topics in Biological, Physiological, and Behavioral Adaptations to Spaceflight. Appendix C
Start Date:	04/15/2019	End Date:	04/14/2023
No. of Post Docs:		No. of PhD Degrees:	
No. of PhD Candidates:	1	No. of Master' Degrees:	
No. of Master's Candidates:		No. of Bachelor's Degrees:	
No. of Bachelor's Candidates:	1	Monitoring Center:	NASA JSC
Contact Monitor:	Whitmire, Alexandra	Contact Phone:	
Contact Email:	alexandra.m.whitmire@nasa.gov		
Flight Program:			
Flight Assignment:	NOTE: End date changed per S. Huppman/HRP and NSSC information (Ed., 3/20/2020) NOTE: End date changed to 3/31/2020 per NSSC information (Ed., 1/22/2020)		
Key Personnel Changes/Previous PI:	Jerri Stephenson is no longer working in the project. James S. Garrett is now a Co-Investigator.		
COI Name (Institution):	Cross, Ernest Ph.D. (NASA Johnson Space Center) Garrett, James (NASA)		
Grant/Contract No.:	80NSSC19K0654		
Performance Goal No.:			
Performance Goal Text:			

	This task is part of the Human Capabilities Assessments for Autonomous Missions (HCAAM) Virtual NASA Specialized Center of Research (VNSCOR). The goal of this research is to develop conversation analysis to measure and mitigate inappropriate trust in virtual assistants. These trust measurements will guide system design, particularly the multimodal interactions and mode switching, as well as how to mitigate over trust and trust recovery. We will use conversation analysis to measure trust at multiple time-scales from real-time interactions to longitudinal monitoring of trust over a long duration exploration mission.		
Task Description:	Conversation analysis provides a promising, but relatively unexplored approach to measuring trust. We propose a conversation analysis at the micro, meso, and macro levels which includes not just the words, but also pauses and facial expressions. Specifically, at the micro-level, conversation elements include voice inflections, pauses between works and keystrokes, gaze shifts, and facial expressions. The meso-level analysis includes words exchanged during interactions with the virtual assistant along with other team interactions as they relate to the automation. At the macro level, conversational analysis considers interaction time, interaction effort, frequency of interaction, turn-taking, barging in tendency, and whether it is the person or the virtual assistant who initiates the interaction. Additionally, prior research into conversational analysis indicates there are novel ways of managing or calibrating trust through the presentation of information, e.g., manipulating the tone and cadence of the system when using speech and through facial expressions (Nass & Brave, 2005; DeSteno et al., 2012).		
	Due to time delays in communication, long duration exploration missions will require greater crew autonomy and greater reliance on automation. For this approach to work trust calibration needs to be engineered into the system. Trust is a critical construct that mediates how well human operators use automated systems, such as virtual assistants, that provide decision support. Trust affects people's willingness to rely on automated systems in situations that have a degree of uncertainty and risk. Trust strongly affects the effectiveness of human-agent collaboration, particularly in the willingness to accept suggestions from a virtual assistant. Knowing whether or not to trust automation can be further complicated by lack of sleep, workload, task risk, and task complexity. Moreover, as we continue to push the limits of intelligent systems and rely on them more as decision aids trust calibration (i.e., operator trust is at a level which matches the automation's capabilities) becomes essential to mission execution.		
	Appropriate calibration of trust requires matching the operator's trust to the virtual assistant's current capabilities. Calibration of trust is not something that can happen once, but must occur throughout the life cycle of the interaction between operator and automated system (Hoffman et al., 2009). Trust is a dynamic construct that continuously increases and decreases due to a number of factors, primary the performance of the automated system, i.e., higher performance leads to higher trust and vice versa. Although much effort focuses on creating more capable and trustworthy automation, less effort has considered the equally important consideration of creating trustable automation. Trustable automation is automation that is understandable and that naturally promotes calibrated trust. Therefore, we aim to create trustable automation by continuously measuring operators' trust unobtrusively and in real-time, and then use this measure to guide the virtual agent to employ one or more countermeasures to calibrate trust and improve human-system performance.		
	References		
	partners in economic exchange. Psychol Sci. 2012 Dec;23(12):1549-56. <u>http://doi.org/</u> ; PubMed <u>PMID: 23129062</u>		
	Hoffman RR, Lee JD, Woods DD, Shadbolt N, Miller J, Bradshaw JM. The dynamics of trust in cyberdomains. IEEE Intelligent Systems. 2009 Nov-Dec;24(6):5-11. <u>https://</u>		
	Nass C, Brave S. Wired for Speech : How Voice Activates and Advances the Human-Computer Relationship. Cambridge, MA: MIT Press, 2005.		
Rationale for HRP Directed Research:			
	The outcomes of this research will make two important contributions to the overall HCAAM VNSCOR effort. First, it will promote more effective interactions and acceptance of virtual assistants. Second, it will provide new analytic techniques for understanding how people work with automated agents as team members. Virtual assistants and other types of agents enabled by artificial intelligence represent an important opportunity to extend human capabilities, but only if they are accepted and trusted appropriately. If people trust the virtual assistant too much they will rely on it in situations that exceed its capability, and if they trust it too little they will fail to engage it when it could benefit the team. One pathway towards appropriate trust is to make the virtual assistant more trustworthy: increase its technical capabilities to accommodate any situation. Another approach is to make it more trustable: communicate its capability and allow its capability to be challenged in its interactions with people. Such trustable technology requires three important advances to the state of knowledge in the field:		
	1. An ability to ascertain how much people currently trust the technology		
Research Impact/Earth Benefits:	2. An ability to convey uncertainty and its capability, particularly as part of conversational interactions		
	3. Interaction affordances that provide the opening for people to assess the capability of the assistant, particularly as part of conversational interactions.		
	These three advances for trustable technology require the development of new analytic techniques for understanding human interaction with automated teammates. Real-time, unobtrusive measures of trust represent a particularly valuable, but challenging measure to develop. Trust is most often measured with ratings and indirectly through people's decision to rely on automation, which are obtrusive not diagnostic. Conversation and text-based interactions offer a promising, but unexplored way to assess trust. Text analysis has a 50-year history in domains as diverse as psycholinguistics and cognitive science, and more recently natural language processing, affective state assessment, and sentiment analysis. Building on the foundation of text analysis makes it possible for this research to immediately contribute to data analysis of previous and future studies of automation-human teaming, and to contribute to the foundation of conversational agent design.		

	Major goals of the project during this phase: Collect and analyze conversational data from controlled microworld experiment Develop a machine learning model that estimates trust from the audio and lexical components of conversational data Develop a cognitive model that simulates the consequence of interactions with an interdependent agent Conversational measures of trust:
Task Progress:	we have analyzed the data from a controlled experiment and created a machine learning model that estimates trust in an agent from the lexical and acoustical features of conversations with that agent. The objective of this study was to estimate trust from conversations using both lexical and acoustic data. As NASA moves to long-duration space exploration operations, the increasing need for cooperation between humans and virtual agents requires real-time trust estimation by virtual agents. Measuring trust through conversation is a novel yet unexplored approach.
	A 2 (reliability) × 2 (cycles) × 3 (events) within-subject study on habitat system maintenance was designed to elicit various levels of trust in a conversational agent. Participants had trust-related conversations with the conversational agent at the end of each decision-making task. To estimate trust, subjective trust ratings were predicted using machine learning models trained on three types of conversational features (i.e., lexical, acoustic, and combined). After training, model inference was performed using variable importance and partial dependence plots. Results showed that a random forest algorithm, trained using the combined lexical and acoustic features, was the highest performing algorithm for predicting trust in the conversational agent [Equation]. The most important predictor variables were a combination of lexical and acoustic cues: average sentiment considering valence shifters and the mean of formants, Mel-frequency cepstral coefficients (MFCC), and standard deviation of fundamental frequency. Precise trust estimation from conversational cues and acoustic cues. These results show the possibility of using conversational data to measure trust, and potentially other dynamic mental states, unobtrusively and dynamically. These results have been submitted for publication in the journal, Human Factors, under the title of: "It's Not Only What You Say, But Also How You Say It: Machine Learning Approach to Estimate Trust from Conversation".
	Computational model of interdependent agents:
	We also developed a computational cognitive model of interdependent agents, where one agent is a person and the other is a conversational agent. Conversational agents are likely to represent automation that has more authority and autonomy than simple automation. Greater authority may lead the agent's goals to diverge from those of the person. Such misaligned goals can be amplified by the situation and strategic interactions, which can further impact the teaming process and performance. These interrelated factors lack a systematic and computational model. To address this gap, we developed a dynamic game theoretical framework simulating the human-artificial intelligence (human-AI) interdependency by integrating the Drift Diffusion Model simulating the goal alignment process.
	A 3 (Situation Structure) × 3 (Strategic Behaviors) × 2 (Initial Goal Alignment) simulation study of human-AI teaming was designed. Results showed that teaming with an altruistic agent in a competitive situation leads to the highest team performance. Moreover, a goal alignment process can dissolve the initial goal conflict. Our study provides a first step of modeling goal alignment and implies a tradeoff between a balanced and cooperative team to guide human-AI teaming design. These results have been accepted for publication at the Human Factors and Ergonomics Society (HFES) conference: "Modeling Goal Alignment in Human-AI Teaming: A Dynamic Game Theory".
Bibliography Type:	Description: (Last Updated: 07/06/2025)
Articles in Other Journals or Periodicals	Li M, Alsaid A, Noejovich SI, Cross EV, Lee JD. "Towards a conversational measure of trust." arXiv preprint server. Posted October 10, 2020. <u>https://doi.org/10.48550/arXiv.2010.04885</u> , Oct-2020
Articles in Peer-reviewed Journals	Chiou EK, Lee JD. "Trusting automation: Designing for responsivity and resilience." Hum Factors. 2021 Apr 27. https://doi.org/10.1177/001872082110099, Apr-2021