

**2020 AFOSR Annual Trust and Influence Portfolio Review
Presentation Abstracts
August 24-28, 2020**

Project Key:

T – Trust

I – Influence

M – MINERVA

S – Science of Science-MINERVA

Alphabetical by Principal Investigator

[S, M] Science Genome: A Scholarly Graph Embedding Framework

PI: Yong Yeol Anh, Indiana University

Information not available

[T] Automation Biases in Human-Robot Trust Interactions

PI: Gene Alarcon, AFRL/711HPW RH

The Automation Biases in Human Robot Trust Interactions is to determine the differences in how humans trust people as compared to how humans trust robots. We have utilized experimental manipulations to determine differences in human-human trust and human-robot trust. It is hypothesized that humans will have higher perceptions of other human's trustworthiness and will perform more trusting behaviors with other humans than with robots. It is also hypothesized that performance will have a greater impact on changes in perceptions of trustworthiness and trusting behavior in human-robot teaming. Interestingly, the research to date has found no trust differences between partnering with a human versus partnering with a robot. We explain these findings in terms of anthropomorphism.

[M, I] Leveraging Wedge Issues Among Populations Susceptible to the Influence or Control of the Islamic State and Prevention of Violent Successors

PI: Scott Atran, Artis International

To identify issues and means that reinforce or undermine appeals of violent extremism we have developed dynamic measures for an experimental platform that can be used online or off to evaluate the relative importance of values, groups, perceptions of spiritual and physical strength, individual and group resilience, social cohesion, and dehumanization. Evidence is presented from Iraq, Palestine, Morocco, and Spain. In addition, we have plotted the evolving network structure behind the 2015-2016 Paris and Brussels ISIS attacks, beginning prior to 9/11, to show what aspects of network structure are most apt to help or hinder successful attacks.

[M, I] Addressing Resilience in the Western Alliance Against Fragmentation: Willingness to Sacrifice and the Spiritual Dimension of Intergroup Cooperation and Conflict

PI: Scott Atran, Pembroke College

We examine how internal political forces and external powers exploit shifting value priorities and group attachments within the Western alliance (e.g., NATO, EU) to fragment and weaken it, and how these same value priorities and group attachments might be best leveraged to bolster resistance to internal fragmentation and key external threats. To illustrate, we present a multi-method approach involving ethnographic fieldwork, lab experiments and online surveys, and machine learning and AI. Evidence is presented from Spain (Catalan Independence Movement), UK (Brexit), and Georgia (NATO partnership, EU-Georgia Association), with a focus on the effects of Russian malign influence campaigns in social media.

[T, I] Moral Agency in Robot-Human Interactions (MARIA)- Perceptions, Trust, and Influence

PI: Jamie Banks, Texas Tech University

The MARIA Project focuses on whether perceptions of a robot's moral agency (abilities to *be* and *do* good or bad) are linked to trust and influence. Building on Year 1's findings (perceptions of minded agency and moral capacity are generally similar for humans and robots), Year 2 focused on interaction-context factors and behavior explanations. Findings indicate both manipulated and perceived interaction types show differing trust outcomes, with playful interactions potentially serving as a 'stepping stone' to promoting trust; additional analysis is underway.

[M] The Social Ecology of Radicalization

PI: Noémie Bouhana, University College London

Information not available

[T] Trust in Human-Robot Interaction via Embodiment and Theory of Mind (THRIVE++)

Prof. Angelo Cangelosi (PI), Wenxuan Mou, Martina Ruocco, University of Manchester, UK

The design of Theory of Mind (ToM) capabilities in interactive robotic agents, which scaffolds on the developmental acquisition of basic social and cognitive capabilities, can allow robots to recognize the actions, goals, intentions, desires and beliefs of others. ToM competence is also key to building and maintaining trust, as it allows the robot to distinguish between reliable (trustworthy) partners and unreliable (untrustworthy) people or situations. This project aims to investigate the complex, dynamic interaction between people's and robots' ToM skills for the building of trust interactions. It specifically aims at the design a computational cognitive developmental architecture for the design of ToM capabilities in robots, and its role in

contextualized human-robot trust interactions, investigating the effects of developmental, social and contextual factors in ToM and trust. Initial human-robot experiments have shown that robots possessing a higher-level of ToM abilities (i.e. attitudinal agency) were trusted more than the robots with lower-level ToM skills (i.e. mechanical agency).

[T] A Longitudinal Study of Trust Calibration Methods with Individual Differences

PI: Fang Chen, University of Technology Sydney

This project aims to achieve a better understanding of human decisions influenced by their trust on machines. Critical factors including machine performance, task categories and subjective perception differences are involved in the investigation. Through this study we will be able to formulate and quantify the trust relationship between human and machine, and lead to better collaboration in a human-machine team.

[T] Improving Situation Awareness in Distributed Human-Robot Teams

PI: Nancy Cooke, Arizona State University

In order to facilitate team situation awareness, trust, and performance in human-robot teams we ask the following research questions: *What are the essential cognitive characteristics of robots required for team situation awareness and effective teaming in distributed human-robot teams? What representations and algorithms are needed to enable robots to realize such behaviors?* This research is accomplished by a multidisciplinary team from human systems engineering and computer science.

[M] Title: Emotions, Attitudes, and the Moral Marketplace: The Psychological, Social, and Situational Determinants of Prosociality and Antisociality

PI: Daniel M.T. Fessler, University of California – Los Angeles

We have demonstrated that people are influenced by the immediate prosocial behavior of others, as observing altruism can elicit *elevation*, an uplifting emotion that mediates increases in prosocial motivation. The propensity to experience elevation is demonstrably a function of prior expectations of others' cooperativeness, itself likely a product of previous experience. This cognitive-affective-social framework can shed light on individual differences, group differences, and context-specific effects in the spread or dissolution of prosociality.

[T] Working with or around you? Investigating team growth and adaptation with a cooperative machine agent

PI: Gregory Funke, AFRL

Strategic research guidance from the U.S. Department of Defense (DoD) and U.S. Air Force states a need for advanced autonomous/agent systems capable of human-machine teaming (HMT) to meet critical Air Force objectives. The goal is to eventually develop sophisticated machine agents that act as full teammates, able to contribute to team planning and capable of executing complex tasks with minimal human oversight.

Present-day agents are largely unable to meet these Department of Defense (DoD) objectives, in part because of limitations in their ability to interpret others' intentions to and communicate effectively. Until fully autonomous agents with these capabilities are realized, it is likely that the DOD will rely on semi-autonomous agent systems to partially fulfill the need for HMT. These agents are likely to have the capability to handle narrow unknowns and only a limited ability to communicate and aid their human teammates in critical aspects of teamwork, such as the development and adaptation of team strategies. Regardless of these limitations, agents will need to act interdependently with teammates in pursuit of common goals, which in turn will hinge on the ability of the agent teammate to cooperate by adapting its tactics to support the strategic direction of the team. The inclusion of semi-autonomous teammates with different capacities to cooperate in this manner is likely to have an impact on critical team processes, such as the formation of shared mental models (functional understanding held in common by teammates), team communications, and team trust, all of which will have consequences for team performance. However, the ways these effects will manifest in the context of HMT is largely unknown. Furthermore, the nature of these processes are likely to evolve as human teammates acquire a greater understanding of an agent's capabilities and limitations.

In response to these gaps, our research approach in this project includes a manipulation of a machine agent's capacity to cooperate (CTC) with its human teammates and a manipulation of team experience. We will examine the effects of these factors on team performance and critical team processes, such as communication, shared mental models, and trust, in a computerized version of the complex and strategically oriented board game *Pandemic*. Our efforts will include a longitudinal design where we will collect an array of behavioral, physiological, and self-reported data from teams of three humans as they play *Pandemic* with an agent teammate over the course of five experimental sessions.

[M, T] The Warfighter's Tolerance for Autonomy Strategy & Systems

PI: Jai Galliot, University of New South Wales

Information not available

[M] The Strength of Social Norms Across Cultures

PI: Michele Gelfand, University of Maryland

Humans are unique among all species in their ability to develop, maintain, and enforce social norms. It is likely, then, that humans have evolved cognitive and affective neuro-mechanisms to be able to detect norm violations very quickly, which affords punishment of violators and enforcement of the social order, though there is very little research on this topic. In this presentation, I will present data from a number of EEG and fMRI studies that have addressed

questions such as: How can we develop new measures of detection of social norm violations at the neurobiological level? Which cultures, individuals, and situations show stronger neurobiological reactions to norm-violating events? How are neurobiological processes underlying social norms violations related to behavioral processes, including attitudes, self-control, creativity, among other behaviors? Do neurobiological indices of social norm violation mediate cultural differences in attitudes and behaviors? I will also discuss how research on the neurobiological basis of social norms can help to improve intercultural interactions.

[M, T, I] Project Title: Organizational Implications of Autonomy-Mediated Interaction

PI: Jonathan Gratch, University of Southern California

Abstract: In human organizations, people often act through intermediates (e.g., a leader may be represented by junior officer or a client represented by a lawyer). There is growing interests in AI algorithms fulfilling this mediating role (e.g., management functions may be delegated to algorithmic management and autonomous agents may negotiate disputes on our behalf). This project draws on computer and organizational science expertise to study how such autonomous representatives shape people's cognitions, behavior and organizational outcomes. We will report on a series of studies that give insight into these processes in such areas as ethical behavior, worker status and employee voice.

[T] Project Title: Development of Analytic Representations of a Computational Cognitive Architecture

PI: Glenn Gunzelmann, AFRL/711HPW RH

Abstract: The overall focus of this research effort is to create methodologies that allow simulation-based models of human cognition to be expressed in analytic forms to promote both advancement in our understanding in human cognition and application of cognitive modeling in complex environments requiring real-time assessments of cognitive performance. With regard to theory, we address longstanding challenges related to model comparison and synthesis when looking across modeling formalisms. For instance, we have identified mathematical relationships between ACT-R's declarative memory system and the Lognormal Race framework, allowing us to derive likelihood functions for model variants of the fan effect. And, we have developed a tutorial showing how to develop ACT-R models with various methods, including fast Fourier transformation, the Lognormal Race frame, discrete time and continuous time Markov processes. These capabilities enable formal model comparison across frameworks, and promotes more deliberate and detailed assessment, comparison, and synthesis of model mechanisms in the understanding of cognitive phenomena. On the applied side, the ability to cast simulation-based models in an analytic framework eliminates the requirement for Monte Carlo simulation to establish model behavior and estimate parameters to fit specific sets of data. This enables real-time parameter estimation based upon streaming data, establishing a foundation for performance and human state monitoring in operational settings. This talk will review research accomplishments associated with this project and highlight the opportunities created through the methodological advancements in both basic and applied settings.

[I] Networked Social Influence and Acceptance

PI: Dr. Weisi Guo, Cranfield University

[T, I] Acoustic-Prosodic and Lexical Cues to Deception and Trust: Deciphering How People Detect Lies

Sarah Ita Levitan, Michelle Levine, and Julia Hirschberg (PI)

Humans rarely perform better than chance at lie detection. To better understand human perception of deception, we created a game framework, LieCatcher, to collect ratings of perceived deception using a large corpus of deceptive and truthful interviews. We discuss the acoustic-prosodic and linguistic characteristics of language trusted and mistrusted by raters and compared these to characteristics of actual truthful and deceptive language and to a classifier we have built to automatically distinguish trusted from mistrusted speech.

[T] Development of a Remote-fNIRS Device for use Under Naturalistic Working Conditions

PI: Leanne Hirshfield, Syracuse University/University of Colorado

Information not available

[T] Project: Toward Undifferentiated Cognitive Agents: Determining Gaps in Comprehension

PI: Pascal Hitzler, Kansas State University (was: Wright State University)

We report on work performed as part of a collaboration between three concurrently funded projects (AFRL: Myers and Marujama; Drexel: Salvucci; Kansas State: Hitzler) which tightly collaborate on developing undifferentiated cognitive agents. The combined project goal is to develop methods and prototypes which enable a cognitive agent to process and deal with natural language instructions for solving given tasks. In particular, in this presentation we will report on the rationale and progress regarding the use of knowledge graphs and ontologies as knowledge bases underlying a cognitive agent. Partner project presentations will report on other aspects of the combined work.

[T] An Experimental Investigation of how Robotic Learning Affects Trust

PI: Nhut Ho, California State University, Northridge

The primary objective of this experiment was to determine if and how an operator's trust in an adaptive system that incorporates machine learning is affected by the presence/absence automation transparency about this machine learning. The experiment simulated an adaptive system which could sometimes allocate decision making in a search and rescue task to the

participant, and sometimes to the automation. While the main prediction that trust would increase with increased transparency was not confirmed, statistically significant main effects, which could be attributed to learning, were found for the different query types (incorrect and correct detection) the operator answered.

[M, T] The Disruptive Effects of Autonomy: Ethics, Trust, and Organizational Decision-making

PI: Michael Horowitz, University of Pennsylvania

Information not available

[T] Project: A neuro-psycho-ergonomic framework for studying long-term bonding between humans and autonomy

PI: Frank Krueger, George Mason University

Trust pervades nearly every social aspect not only of human-human but also human-autonomy interactions. Although trust plays a crucial component in developing long-term bonding between humans and autonomy; however, its neuro-psycho-ergonomic underpinnings are not well understood. This project will present an overview of scientific approaches --ranging from psychology, ergonomics, and neuroscience-- to develop a unified framework in studying the neuro-psycho-ergonomic mechanisms of long-term bonding between humans and autonomy.

[M, I] Mobilizing Media: A Deep and Comparative Analysis of Magazines, Music, and Videos in the Context of Terrorism

Anthony F. Lemieux, PI (Georgia State University); Carol Winkler, co-PI (Georgia State University); Jonathan Pieslak, co-PI (City College of New York); Akil Awan, co-PI (University of London); Weeda Mehran, co-PI (Exeter University, previous Post-Doctoral fellow at GSU); Ben Miller (Emory University); Nelly Lahoud (Consultant); Jarret Brachman (Consultant); Humera Khan (Consultant).

Abstract / Summary / Description:

Mobilizing Media utilizes three approaches to analyze the media campaigns of predominantly MENA-based extremist groups. The first examines audience cultivation through identification of core themes and persuasive elements, their level and nature of use by different groups, and the events that prompt them to change over time. The second examines the unique information conveyed by textual, visual, aural, and sonic elements of the media campaigns, as well as strategies of interactive reinforcement the groups deploy across the different modalities. The third examines how the groups' use of information, motivational appeals, and behavioral skills maximizes their chances of producing behavioral change. This presentation will summarize key findings over the life of the project to date, and will outline future directions.

[S, M] Dynamics of the Scientific Ecosystem

PI: Daniel Larremore, University of Colorado

The timing and pace of discoveries are impacted by the dynamics of scientific careers, including the noisy sorting process of the job market and the impacts of parenthood at various career stages. Here we report updated findings on the dynamics of the scientific ecosystem through the lens of individuals' careers, through studies of the assembly and stability of prestige hierarchies in faculty hiring, the unequal impact of parenthood across academic research, and the role of graduate student researchers in explaining the outsized productivity and impact of the scientific elite.

[I] Project Title: Digital Deception: The Cognitive and Social Mechanisms of the Spread of Fake News

PI: Kristina Lerman, USC Information Sciences Institute

The spread of disinformation in online social networks highlights our vulnerability to digital influence. The goal of this effort is to elucidate the social and cognitive processes involved in the spread of online misinformation and how to mitigate its adverse effects. Our research program has combined theories from cognitive science with network analysis to quantify social, network and psycho-linguistic factors involved in the spread of deceptive information. We analyzed millions of Twitter accounts that interacted with Russian Internet Research Agency (IRA) trolls to identify traits predictive of who promotes information trolls share, and characterize deceptive language used by these accounts. In ongoing work, we examined how structure biases perceptions of popularity in online social networks. Understanding how hostile actors manipulate attention and perceptions of susceptible individuals will lead to new approaches for reducing their vulnerability, by steering attention to more desirable content, or by altering their network structure to reduce hostile influences.

[M] Neural Bases of Persuasion and Social Influence in the U.S. and the Middle East: A Social Neuroscience Approach to Ideological Interpersonal Conflicts

PI: Matthew Lieberman, University of California – Los Angeles

The history of experimental persuasion research is focused on topics that targets are not already ideologically 'dug in'. While conversion events do occasionally cause dramatic shifts in ideological positions, there are no known brief psychological interventions that reliably produce these shifts. We propose that there is less value in focusing on attempting to shift ideological views than in trying to shift how people view and interact with people who have different ideological views from themselves. I will present data focused on three interrelated questions of this research program: (1) Can we use functional near infrared spectroscopy (fNIRS) with machine learning to predict an individual's ideological position from simply watching a video? (2) Can we use fNIRS to identify when ideological interpersonal conflict is occurring during a realtime naturalistic conversation? and (3) Can we use interventions such as self-affirmation to increase open-mindedness to opposing ideological views, not as a position one should hold personally but as a position one can understand another person holding. The portability and

mobility of fNIRS holds promise for examining real world social interactions, like ideologically heated conversations, and identifying psychological processes occurring in the moment (that could not be reported on without undermining the interaction) that reflect further entrenchment of one's position, disdain for the interaction partner, or alternatively, increasing open-mindedness to others holding positions different from one's own.

[M, I] Title: The Dynamics of Common Knowledge on Social Networks: An Experimental Approach

PI: Read Montague, Virginia Tech, and Gizem Korkmaz, University of Virginia

We study protest as a collective action problem where an individual wants to participate only if joined by "enough" like-minded others. In game-theoretic contexts, coordination requires that agents know about each other's willingness to participate and that this information is common knowledge among a sufficient number of people. Stylized models combine social structure and individual incentives together, and provide a rigorous game-theoretic formalization of common knowledge, and the characterizing network structures. These models emphasize simple node-to-node or bilateral communication, or study the effects of "richer" on-line communication mechanisms, such as Facebook. However, very little is known about how well these models explain phenomena in practice; nor do these models incorporate individual and behavioral factors, and psychological processes. We aim to empirically test novel hypotheses at both individual and group levels by conducting controlled human subjects experiments in three different environments (laboratory, online, and neuro-imaging) that will inform models of collective behavior. The objectives of this integrated framework are (i) to characterize how different communication mechanisms can facilitate actionable common knowledge through local interactions, (ii) to understand the effect of network structure, and (iii) to quantify individual and group level behaviors, and neural processes that affect its formation.

[T] Project: Towards Undifferentiated Cognitive Agents for Diverse Specializations

PI: Dr Christopher Myers (AFRL/711HPW RH)

CoPI: Dr Benji Maruyama (AFRL/RX)

CoPI: Dr Christopher Stevens (AFRL/711HPW RH)

We report on research that is part of a three-team collaboration between AFRL, Kansas State University, and Drexel University toward the development of *undifferentiated cognitive agents*: computational cognitive models capable of being taught through instruction and interaction. We predict that undifferentiated cognitive agents will reduce modeling costs relative to current approaches. We will provide an overview of the goals and approach taken to achieving undifferentiated cognitive models/agents and provide a cost accounting of current modeling approaches in our targeted development environment. We currently have a working end-to-end system that we are striving to scale in its ability to learn increasingly complex tasks.

[M] The Social and Neurological Construction of Martyrdom

PI: Robert Pape, University of Chicago

Co-PI: Jean Decety, University of Chicago

Project Point of Contact: Jay Collins, University of Chicago, fjc@uchicago.edu

This project uses the tools of social science and neuroscience to understand the role of propaganda used by violent extremist organizations such as ISIS to recruit, and the social and neurological factors that explain how and when propaganda motivates individuals in the West to sacrifice for a group and its cause. The core theoretical finding is that ISIS and other Islamist militant groups use two narratives – a social and a heroic martyr narrative – to attract two pools in the West: those with strong ties to Islam and established Muslim communities and those, like converts, with weak ties. To investigate the recruitment appeal of these narratives, we 1) collected and catalogued all known official propaganda videos produced by ISIS, AQC, AQAP, and al-Shabab through 2019, and JN through 2018, more than 6,000 videos, all coded for over 30 visual and thematic elements; 2) systematically analyzed narratives in all ISIS and AQC videos targeted at Western audiences; 3) created a database of demographic, biographic, and motivational variables of all US jihadist offenders between 2001 and 2019 (N=559), including variables on the role of propaganda in mobilization; 4) conducted a behavioral study (N=283 adults) in which subjects rated social and heroic narrative ISIS clips; and 5) completed an EEG study (N=76) to probe neural pathways which distinguish between narrative types and dispositional factors which predict susceptibility to specific narratives. Findings support the importance of the two narratives and the mechanism of narrative transportation in shaping recruitment appeal of ISIS recruitment videos in Western audiences.

[T] Trust Affordances in Human-Robot Teaming

PI: Dr. Laurel Riek, Computer Science and Engineering, UC San Diego

When engaging in human-robot teaming in uncertain, dynamic environments, it is crucial that there is mutual understanding and well-calibrated trust between humans and machines. This talk will discuss several recent efforts in my lab which address this issue, including: 1) Recently published study exploring trust calibration and error recovery during proximate, mobile co-manipulation tasks, 2) New methods to dynamically measure teaming to inform autonomous robot adaptation and learning, and 3) New approaches for avoiding operational failures in critical environments. I will also discuss our recent efforts in applying this basic research to building and deploying new, low-cost telemedical robots in emergency medical environments to support improved healthcare safety during the pandemic.

[I] Assessing the Impact of Exogenous Shocks on User Behavior and Information Diffusion in Social Media

PI: Daniel Romero, University of Michigan

This project aims to study how exogenous shocks such as political movements, natural disasters, and financial and health crises impact information diffusion dynamics on social media. Our approach combines modeling, large-scale data analysis, and simulations that draw on social science theory and computational methods to understand various information diffusion components such as network structure, user behavior, and linguistic features. Our results will enable us to better predict online information cascades during times of uncertainty and enable the design of effective social media tools based on principled predictions of behavior.

[T] Toward Undifferentiated Cognitive Agents: Translating Instructions to Knowledge

PI: Dario Salvucci, Drexel University

While a number of research efforts have successfully modeled the cognitive process of instruction following, our computational models remain brittle with respect to the given instructions and have difficulty adapting dynamically to variations in the instructions. Our project aims to develop novel models of instruction following applied across a variety of tasks. In this presentation, we will highlight new work to make instruction following more robust, including (1) more flexible grounding of language to execution, (2) processing of instructions that allows for inference of implicit instruction knowledge, and (3) dynamic, interactive clarification of instructions during both the learning and execution stages.

[T] Supporting Trust Calibration and Attention Management in Human-Machine Teams

PI: Nadine Sarter

Graduate Students: Kevin Lieberman and Karanvir Panesar

Organization: University of Michigan

The goal of this research project is to better understand the role of and relationship between trust and attention management in human-machine teams. Specifically, we examine how the interaction of these two phenomena evolves unaided over prolonged periods of time and how it can be shaped through design and training in the interest of improved joint system performance. Candidate approaches to training and mental model development as well as interface designs that support system transparency and attention guidance are being developed and assessed in the context of a multi-UAV control simulation, using eyetracking and traditional trust measures such as subjective ratings and performance.

[T] Title: Enabling Trusted Human-Like Artificial Teammates

PI: Matthias Scheutz, Tufts University

This multi-disciplinary project aims to lay the foundations for future autonomous systems in mixed-initiative human-machine teams to be able to operate at human-like levels of interactivity and effectiveness. We are integrating measurements of neurophysiological signals from human teammates (fNIRS, EEG, eye gaze) with multiple additional measurements (physiological, linguistic, behavioral) and situational contextual information to classify various individual and team cognitive states.

Classified cognitive states of all teammates are individually tracked and fused in real-time, and integrated into a shared mental model (SMM) which uses advanced probabilistic “theory of mind” representations to capture team and task states with their associated uncertainties, and supports the decision-making and behavior adaptations of the autonomous artificial teammates.

[T] Title: Trustworthy Human Interaction with Robotic Swarms: FA9550-15-1-0442

Primary Investigator: *Katia Sycara, Carnegie Mellon University, katia@cs.cmu.edu*

Co-Investigators: *Michael Lewis, U. of Pittsburgh; Nilanjan Chakraborty, Stony Brook*

We are studying trust calibration in human supervisory control of a remotely situated, autonomously coordinating robotic swarm of Uninhabited Vehicles. Some research questions we are asking are: How can we robustly measure human trust in a swarm as a mission evolves? Since trust is a subjective attitude, how can we construct quantitative predictive trust models of humans that capture dynamic changes in human trust? In this briefing we present experimental studies of dynamic changes in human trust in a swarm as a mission evolves, a computational predictive model of dynamic trust and trust effects on the swarm behavior. Recent work investigates distributed algorithms for swarm self-repair and their effects on both performance and human trust. We have also started to investigate deception in the swarm human context.

[I] From Online Actions to Attitudes

PIs: Lexing Xie, Australian National University, and Yu-Ru Lin, University of Pittsburgh

The PITT and ANU team curated a set of Twitter and YouTube activities for several online movements (e.g., black lives matter, gun control, abortion, and climate change). We measure the difference of engagement patterns for users with distinct political leaning. We will describe a few research questions that the analysis can help answer. We will also summarise recent work in both labs, such as quantifying the effect of twitter data sampling and linking protests to changes in online prejudice.

[T] Developing human-machine systems that actively calibrate a user’s trust

PI: Alan Wagner, Pennsylvania State University

Our recent work exploring human-robot trust focuses on several different problems. We consider trust during an evacuation, the impact of explanations on trust, how trust degrades over time, how a robot might recognize signs of trust, and risk-based planning for autonomous vehicles. Our

work, therefore, seeks to create trust-related tools for a variety of different autonomous system applications.

[S, M] Fundamental Dynamics, Predictability, and Uncertainty of Scientific Discovery

PI: Dashun Wang, Northwestern University

Information not available