HARVARD | BUSINESS | SCHOOL

Rising Scholars Conference
Technology and Operations Student Research Presentations

Nicholas Brown **Virginia Tech**
nichb15@vt.edu

Nicholas Brown has the mindset of a lifelong learner and a voracious appetite for reading. These attributes have helped him to earn four Master's degrees, with each degree informing his privacy and cybersecurity research in the Information Systems discipline. Notably, his favorite degree is the MA in Christian Leadership, where he learned the principles of servant leadership from The New Testament and practices it daily with his students, colleagues, family, and friends.

ABSTRACT:

The privacy-explainability paradox in machine learning systems:
An empirical examination of the role of human involvement in training data privacy

Machine learning (ML) systems are data-driven statistical algorithms that perform pattern-matching functions to draw inferences from large collections of data, often referred to as training data. Inherent in ML systems is a dichotomy between privacy and explainability called the privacy-explainability paradox.To explain any decision or output rendered by an ML system, it almost always requires viewing the training data, whether to evaluate the data for possible evidence of biases, fairness, or accuracy. However,training data may include data subjects' biometrics that can be used to identify their natural persons, a plausible privacy violation. Extant human-ML research recognizes the importance of human involvementto augment recommendations rendered by algorithms. Yet, when viewing human involvement from a privacy lens, individuals have express concerns when human agents view and annotate users' biometric data. This paper investigates the role of human involvement versus machine (technical) processing of voice training data in the context of virtual assistant use (e.g., Amazon Alexa, Alibaba AliGenie, Apple Siri, and the like), and users' willingness to share their biometric voice data with companies to train a company's voice-based algorithms. We operationalize the construct of *perceived human involvement* through the use of privacy policies, used to describe a company's data handling practices and logic, as theform of explanation presented to users, and draw upon the antecedents-privacy concerns-outcome (APCO) model to investigate users' privacy concerns. We employ a scenarios-based survey using a between-subjects 24 factorial design with *perceived human involvement* (human vs machine processing of data), *virtual assistant type* (Amazon Alexa vs Alibaba AliGenie), *cloud provider* (Amazon Web Servicesvs Alibaba Cloud), and *level of information processing* (high vs low) to investigate users' privacy concerns and trust toward a company and their subsequent *willingness to share their voice data* for ML training purposes.

Based on pilot study findings (n=76), we observe that users in the human involvement condition with privacy policies expressing details about the role of human annotators in the training of ML systems, namely through supervised learning and data annotation practices, show higher privacy concerns and higher privacy risk perceptions than users in conditions that receive privacy policies mentioning only machine processing of users' data (no reference to human involvement). Further, we find that users in the human involvement condition show lower willingness to share their voice training data than users in the machine processing condition, revealing preliminary evidence of the privacy-explainability paradox, whereby greater transparency about the role of human involvement in ML augmentation leads to users' increased privacy concerns and decreased willingness to share training data. We further discover findings contrary to our initial expectations. Users in the Amazon Alexa condition show lower levels of trust and higher levels of distrust than users in the Alibaba AliGenie condition, even though users were unfamiliar with the Alibaba AliGenie virtual assistant. This implies that a possible *default distrust bias* exists in users toward companies they know, to the extent users are more willing to trust and share their data with acompany and product they have not used before. We intend to launch a full data collection for the follow-up study to fully explore the roles of human involvement and default trust bias on virtual assistant users' privacy decision making and behaviors.

Using a privacy lens, we extend discussion in the human-ML augmentation discourse by conceptualizing a new construct, *perceived human involvement*, as a proposed antecedent to the APCO model, one of which can predict outcomes related to individuals' willingness to share training data and to opt-out of training data participation. These implications stimulate rethinking toward the privacy-explainability paradox. Moreover, for researchers and practitioners investigating human-ML dynamics, privacy concerns should be strongly considered, because new legislation may empower consumers to restrict the sharing of their training data, thus inhibiting the sufficient development of ML systems and their ability toobtain and

process data sets with complete and diverse features, necessary to reduce algorithmic biases and systematic unfairness in outputs rendered by these systems.

Dustin Cole **Michigan State University**
coledust@msu.edu

I am a 5th year PhD candidate at Michigan State University in the Supply Chain Management Department. My research focuses on how organizations can have a beneficial impact on the communities in which they operate. I examine these topics through both a social and environmental sustainability lens, exploring topics of inclusive employment and water respectively.

ABSTRACT:

Does leader disability status improve performance outcomes for workers with a disability?
An empirical study in the apparel industry.

This research, part of my dissertation, examines the impact of leader-worker disability status similarity on front line manufacturing worker productivity using micro-data gathered from a real-world organization, building on the nascent field of inclusive operations and exploring how organizations can both be profitable as well as inclusive of traditionally marginalized workers. The essay examines the moderating influence of direct supervisors with a disability on workgroup productivity as the number of workers with disabilities increases. Results show that a direct supervisor with a disability does indeed benefit the productivity of workers with disabilities. However, this benefit is not from a direct improvement to productivity, but rather a mitigation of potential productivity declines as the number of workers with a disability increases in the workgroup. The mechanisms are further explored in a case study based on interviews with workers and supervisors from two different organizations that employ a large proportion of workers with a disability. The interviews reveal that workers with a disability place an emphasis on stability in task assignment due to greater challenges in learning, something supervisors with a disability seemed more sensitive to, while the supervisors without a disability focused on more traditional workgroup strategies like cross-training and task flexibility.

Schenita Floyd **University of North Texas**
schenitafloyd@my.unt.edu

Schenita Floyd is a recent graduate of the University of North Texas in Denton, where she obtained her doctorate in Interdisciplinary Information Science with a concentration in data science. Her research interests include Artificial Intelligence (AI), team collaboration, digital workforce transformation, and the gender pay gap. Schenita has a BS degree in Electrical Engineering from Texas A&M University in College Station, and she has an MBA in Finance from Southeastern University.

ABSTRACT:

Artificial Intelligence Teammates in A Collaborative Information Seeking Environment from The Perspective of Women Engineers in The United States

Artificial Intelligence is a disrupter, an enhancer, and an integral part of the workplace. Manyworkplace technologies incorporate artificial intelligence, whether known, unknown or hidden from employees in the form of tools. Collaboration tools such as Microsoft Teams, Slack, and Trello have artificial intelligence features that operate as simple task bots assisting with information seeking. However, they have the potential to become an artificial general intelligence teammate bot in the future. The purpose of this study was to collect design requirements from women engineers on artificial intelligence teammates. A mixed methods research design was used for this study with an online survey and semi-structured interviews to answer the following research questions:

- What is the relationship between self-efficacy and women engineers' attitudes towards AI?
- How do women engineers collaborate in the workplace with or without AI?
- What AI teammates design requirements and features do women engineers recommend?

The study results revealed design requirements from women engineers, including solutions to sociotechnical issues that could arise from artificial intelligence teammates in the workplace. The results showed various ways women engineers

collaborate in the workplace with and without artificial intelligence. Additionally, women engineers' attitude towards artificial intelligence was examined to identify a correlation to self-efficacy. This research study fills a previous study gap that solicited design requirements from research scientists by soliciting practitioners. Practitioners such as women engineers are underrepresented in the workplace, and they could benefit from an artificial intelligence teammate with their design requirements.Additionally, the implications of this study add to the research agenda and improve the knowledge of practitioners in the workplace. Current and future designs of AI have implications from this study. Solutions provided by the participants can be implemented today and in the future. All industries can use preventable measures discussed in the women engineers' discourses. Silent voices can be heard to ensure inclusion in the workplace and the research community. Finally, this study contributes to the information science literature on collaborative information seeking, artificial intelligence design, and engineers' information seeking behaviors.

Keywords: Artificial Intelligence and work, team collaboration tools, collaborative information seeking, women engineers, human-computer interaction, human-machine teams

Jonathan Gomez Martinez **Emory University**
jgomezm@emory.edu

Jonathan Gomez Martinez is the son of Mexican immigrants to the United States and a 3rd PhD student at the Goizueta Business School of Emory University. His research interests are in platform strategy and governance - specifically how platforms can better manage the communities which they facilitate. This is his first academic presentation outside of Emory, and he is excited to share this experience with you all.

ABSTRACT:

Planned Obsolescence Strategies of Desktop OS Platforms:
An Empirical Investigation of Microsoft-Apple Duopoly Through Developer Versioning
Ramnath K. Chellappa, Jonathan Gomez Martinez, Anand Swaminathan

The desktop operating systems (OS) market is characterized by a duopoly ofMicrosoft's Windows and Apple's MacOS, together accounting for over 80% ofthe desktop OS market since 2009. Both OS platforms regularly put out newreleases with improved features over older release. Both platforms also regularlyannounce the end-of-life (EOL) of prior releases, marking the end of additionsand improvements to the OS release and signalling to the market that the EOLOS is outdated. However, 22% of Windows 7 users continued to be content withtheir outdated OS release 6 years after EOL (Faro 2021) while 1 in 8 Macs ranan unsupported macOS version (Keizer 2016).

The problem of eliciting repeat sales from consumers has been explored ex-tensively in economics as a practice termed "planned obsolescence". Planned obsolescence, as defined by Bulow (1986, p.729), is the "production of goodswith uneconomically short useful lives so that customers will have to make re-peat purchases". As such planned obsolescence strategies have been developedfor durable goods (Choi 1994; I. H. Lee and J. Lee 1998). However, with theexception of Iizuka (2007) and a handful of others, work on planned obsoles-cence has almost exclusively relied on analytical models with little empirical ev-idence. Additionally, prior work examining planned obsolescence has provideda product-oriented explanation rather than a platform-based one - where thevalue to a user is a direct function of the number of other users and contributors(Parker and Van Alstyne 2005).

Taking a platform-centric view of planned obsolescence, we propose an oper-ationalizeable definition of success for a platform's planned obsolescence strat-egy as being minimally dependent on market forces outside of the firm and,accordingly, allowing for frequent introductions of product releases. Further, weconstruct and estimate user and developer utility models combining methodspreviously used to quantify network effects (Chu and Manchanda 2016; Gandal,Kende, and Rob 2000; Zhou, Zhang, and Van Alstyne 2020). Recognizing thatthe developer side of the desktop OS market is itself an industry, we apply atechnology adoption framework (Saloner and Shepard 1995) and control for firmcharacteristics recognized to be influential in driving technological adoption indurable goods markets (Hannan, Pólos, and Carroll 2007).

We rely on several data sources from the platforms and other vendors toconstruct weekly panels for each OS release's adoption and abandonment peri-ods. Using a collection of Arellano-Bond (Arellano and Bond 1991) and hazardrate (Shankar and Bayus 2003) models, we identify the drivers of adoption andsupport by users and developers, respectively, as well as abandonment by both.With estimates of the decision-making process of both users and developers, weare able to discuss the success of both platforms' planned obsolescence strategy.We find that MacOS relies less on platform mechanisms to promote new re-leases than Windows does, and this allows for a more frequent release schedule for MacOS than for Windows.