Rising Scholars Conference
Information Systems & Operations Student Research Presentations

Cassidy Buhler
cb3452@drexel.edu

Drexel University
LeBow College of Business

Cassie Buhler is a Ph.D. Candidate in Business Analytics at Drexel University with a graduate minor in Computational Data Science. Prior to this, she received a B.S. in Mathematics from the University of Utah. Her research is at the boundary of optimization and machine learning, focusing on computational methods for mixed-integer and derivative-free optimization. She is interested in using these tools to tackle real-world problems, specifically for conservation and natural resource management decisions.

Abstract:

Decision-Making for Land Conservation: A Derivative-Free Optimization Framework with Nonlinear Inputs

Authors: Cassidy Buhler and Hande Y. Benson, Drexel University

Protected areas (PAs) are designated spaces where human activities are restricted to preserve critical habitats. Decision-makers are challenged with balancing a trade-off of financial feasibility with ecological benefit when establishing PAs. Given the long-term ramifications of these decisions and the constantly shifting environment, it is crucial that PAs are carefully selected with long-term viability in mind. Operation research tools such as simulation and optimization are commonly used for designating PAs, but current decision models are primarily linear and consequently, have limited applications.

To address this gap in the literature, we proposed a derivative-free optimization framework formulated as a mixed integer nonlinear programming (MINLP) problem. This framework enables decision-makers to consider more complex scenarios that are not adequately captured by linearity. As a proof of concept, we used a nonlinear component, population viability analysis (PVA), to incorporate the risk of extinction in our framework. That is, our framework selects the optimal sites of PAs which maximize the probability of survival for a given species. Our numerical results show our models yield PAs with similar expected risk to that of preserving every parcel in a habitat, but at a significantly lower cost.

This research promotes interdisciplinary work by providing a new mathematical programming tool for conservationists that allows for linear and nonlinear inputs, continuous and discrete variables, and can be paired with existing ecological software.

Xufei Liu
xufei@wharton.upenn.edu

University of Pennsylvania
The Wharton School

Xufei Liu is a second-year Operations Management Ph.D. student from Wharton's OIDD department. Her current interests include finding methods for analyzing and modeling high-dimensional, complex datasets and designing intervention systems for performance improvement. Other interests include novel applications of generative AI to new dataset forms, especially in health-related fields.

Abstract:

Elevating Operations: A Sensor-Powered Performance Leap
Authors: Xufei Liu, Gad Allon, and Ken Moon; The Wharton School, University of Pennsylvania

For military aviators, functioning within a high-stress, no-fail work environment is a given. They have to make split-second decisions in the air, piloting 44,000 lbs of machinery while experiencing up to 9G's of acceleration. The learning curve is steep and potentially deadly - aviators may be affected by hypoxia under these intense stressors, which can lead to loss of consciousness and death in severe cases. Increased fatigue may also compromise their cognitive performance and result in cutting training missions short - an expensive consequence due to the high costs associated with military training flights.

We outfit aviators with multiple sensors that track environmental conditions and biological responses, ranging from linear accelerations to heart rate. Through these sensors, we attempt to detect the onset of aviator fatigue and increased stress so that we may intervene to mitigate performance deterioration. as there are currently very few measures in place to monitor and track fatigue. Hence, our overarching question is: How can we detect optimal points of intervention mid-flight? We have found evidence of aviator performance improvement (both physically and mentally) as they gain flight experience. While there is no replacement for experience and flight time, we hope that proper intervention and suggestions can enhance the rate at which aviators improve their performance - this can increase their safety and help them get the most out of each training flight.

However, before we can find points of intervention, we must first be able to decipher the current state of a pilot's performance and stress level given just the raw, high-dimensional sensor data. This data is often noisy and may have missing points since sensors do not function perfectly during field experiments. Thus, while traditional operations management problems often have a defined state and action space, the complexity of our dataset means the state space of our problem is not clearly defined. Thus, our primary focus of this talk is: How can we gain insight when working with a high-dimensional sensor dataset when the raw readings are too complex to understand?

We create two methods of distilling the high-dimensional data into an interpretable,

actionable model. First, we can statically analyze the entire flight trajectory using Convolutional Neural Nets, and with gradient tracing, we can create class activation mappings to determine points of a flight with good and bad performance. From this model, we find that there exists a temporal element to feature importance within our data - different features of our sensor measurements play varying roles of importance in determining overall performance level of a

pilot, and changes as time changes. Second, we can create a deep state-space model to find latent Markov states that are lower dimensional to understand the evolution of the model over time.

The appearance of complex, noisy data will become more common as technology progresses and sensors grow more affordable. It has already emerged in healthcare where nurses wear sensors to track physical stressors, in rental car services to distinguish reckless vs safe driving, and in dangerous work environments to track the well-being of firefighters. Our goal is to create a new method for dealing with these forms of data so that we can simplify the problem to a sufficient lower-dimensional latent representation.

Zanele Munyikwa
zanele@mit.edu

MIT
Sloan School of Management

Zanele Munyikwa is a PhD candidate in the Information Technologies group at the MIT Sloan School of Management, where she has been recognized as an MIT Presidential Fellow and is supported by the MIT Sloan Doctoral Fellowship and the Accenture Convergence Fellowship. Her current research focuses on the economics of digitization with a focus on the impact of artificial intelligence on labor, work and employment. Her primary methodologies are online experiments and machine learning for causal inference in observational data. Prior to starting her PhD, Zanele was a Research Fellow at Stanford Graduate School of Business. She earned a BS in Computer Science with honors at Duke University. She has an SM in management research from MIT.

Abstract:

Machine Learning for Hedonic Wage Analysis: Estimating the Value of Writing in the Digital Age

The rapid development and enhancement of generative artificial intelligence models tailored for writing have brought increased attention to the economic significance of writing skills. This shift aligns with broader discussions initiated by Deming and Kahn (2018) and others, highlighting the transformation of the labor market due to computerization, where the emphasis is increasingly placed on non-routine tasks, and compelling evidence highlighting the rising value of social skills. In this paper, I collect a rich set of job postings data and extract features using natural language processing methods to estimate the value of writing as a skill in the U.S. labor market. Traditionally, researchers have predominantly relied on hedonic wage regression to measure skill premia, as demonstrated by Autor and Handel (2013). However, these approaches are constrained by their linear assumptions and susceptibility to omitted variable bias.

Fortunately, the availability of big data allows for better controls in estimating implicit prices in these models. We leverage employment data from the Bureau of Labor Statistics, occupational data from O*NET, and over 5 million detailed job postings and salaries from Greenwich.HR. We apply machine learning techniques, including lasso, double selection, regression tree models, gradient boosting, and random forest models, to estimate the returns on writing skills. Furthermore, we engage in a comparative analysis, juxtaposing our estimates with those derived from a series of text injection experiments inspired by Bana (2021). To enhance our approach, we fine-tune an open-sourced large language model, Llama2, employing a dataset comprising 5000 job postings paired with their corresponding salary information. Our initial objective is to validate the text injection method for estimating skill returns, with a specific focus on writing as a skill, by benchmarking these estimates against the machine learning method that exhibits the most robust goodness of fit, specifically XGBoost. Subsequently, we extend this investigation to include additional text injection experiments that yield more nuanced, skill-specific insights. This expanded analysis offers estimations of how various writing skills, such as technical writing, content creation, and copywriting, are valued within the context of job postings.

Finally, our comparative analysis over time reveals a discernible and substantial growth in the perceived value of writing as a skill over the past five years. This study contributes valuable insights into the dynamic interplay between technology, skills, and the evolving landscape of

employment, shedding light on the ever-increasing significance of effective written communication in the modern workforce.

Renzhi (Fred) Zhao
rzhao3@gsu.edu

Georgia State University
J. Mack Robinson College of Business

I am a fifth-year Ph.D. Candidate and Researcher in the Computer Information Systems Department at Georgia State University. My scholarly pursuits are primarily rooted in exploring the human decision-making process, open-source software development and the nuanced interactions between humans and artificial intelligence. Drawing upon multidisciplinary theories from information systems, management and psychology, I aim to contribute to the adoption of AI systems (especially decision-aided AI systems) to everyone in the society. My methodological expertise lies in experiment design and quantitative data analysis. I have amassed a portfolio of academic contributions that are targeted at esteemed business research journals such as Information Systems Research and Management Science.

Abstract:

Experts Want More Details? The Impact of AI-Generated Explanation Completeness on Users' AI-Aided Decision-making Performance among Experts and Novices

Artificial intelligence systems are increasingly being used to help humans make better decisions. AI systems can provide recommendations that are tailored to users by analyzing large amounts of data and identifying patterns that humans may miss. All users with different levels of domain expertise can benefit from using AI systems to aid their decision-making. Experts can use AI systems to augment their insights and domain knowledge, while novice users can use AI systems to make better decisions with limited knowledge/experience and learn decision-related domain expertise.

Deciding how much to rely on increasingly capable AI systems presents a major challenge for users, especially in high-stakes domains like financial investing. Without a proper understanding of the AI systems' reasoning, users may overuse or disuse the AI systems. From theoretical perspectives of mental model theories and IS delegation framework, I find that users' inappropriate AI use (i.e., users transfer their decision-making rights to the AI systems to achieve certain goals) stems from their inaccurate mental models of how and why AI systems recommend certain actions. Building accurate mental models about the reasons behind AI's recommendations is essential for users to increase long-term AI-aided decision-making performance. Providing AI-generated explanations is a common way to help users build more accurate mental models of AI systems. Due to different levels of domain expertise, experts and novices have different needs for AI-provided information. Experts need not only key facts, but also peripheral information that supports key facts to reinforce their current mental models of AI's reasoning; While novices do not need much peripheral information to avoid losing focus on key facts when they develop their unmatured mental model. However, there is no one-size-fits-all solution for AI to provide explanations to users with different levels of domain expertise. Therefore, this study aims to find out the appropriate levels of AI-generated explanation completeness, that optimize experts' and novices' AI-aided decision-making performance correspondingly.

Methods and Expected Results

To test my hypotheses, I’m conducting an experiment in a financial investment context. I'm collecting data via online questionnaires on Qualtrics from experiment participants hired on Amazon Mechanical Turk. There are 3(i.e., no AI explanations, AI explanations with key facts, AI explanations with both key facts and supporting peripheral information) * 2(i.e., Experts and Novices) experiment groups. I look forward to showing my preliminary results and full experiment details if I have the opportunity to present my paper at the Rising Scholars Conference.

Expected Implications

This work will extend our current understanding of the AI-aided decision-making process and AI explanation interface design. It will also contribute to the mental model development theories and IS artifacts delegation framework. It will offer guidance to the adoption of decision-aid AI system products among users with different levels of domain expertise, both experts and novices.