WORKSHOP ON AI IN INFORMATION RESEARCH AND PRACTICE: FOSTERING INTERCONNECTED COMMUNITIES

ASIST SIG AI Workshop 2021
Salt Lake City, UT, USA and Hybrid
October 30, 2021 8 AM – 12 PM Mountain Time
Workshop Website: https://asistaiworkshop.web.illinois.edu/

ORGANIZERS:
Soo Young Rieh, University of Texas at Austin, USA (rieh@ischool.utexas.edu)
Clara M. Chu, University of Illinois at Urbana-Champaign, USA (cmchu@illinois.edu)
Dania Bilal, University of Tennessee-Knoxville, USA (dania@utk.edu)

Accepted Presentations

Artificial Intelligence (AI) practical implications in library operations technical and user services: Pakistani Perspective
Muhammad Yousuf Ali (The Aga Khan University, Karachi Pakistan)

The purpose this study discuss about the Artificial Intelligence (AI) effects on the library technical and users services in the field of library and information science. This paper is an attempt to bring AI possible use in the libraries. Further disclose the current usage of AI Tools in the libraries of Sindh Province. The case study methodology is used to carry out this research and researcher visited the library where existing AI application uses interviewed and personal observations added in the discussion section. The Research results shows that Pakistani libraries started AI Technologies slowly, facial recognition, RFID, Self-check in check out and pattern recognition i.e, Thumb Impression, QR code and Bar codes are used in different types of libraries. Most of AI Technologies are used in special and university libraries. Public and other types of library need to work out on AI Technology its implications.

AI and the creation of artificial cultures
Martha Alvarado Anderson (University of Arkansas at Fayetteville)

This presentation aims at highlighting some topics for consideration when creating AI applications and creating communities of practice. Several technological trends in libraries include a higher usage and more development of drones, computer gaming, digital assistants, virtual learning, and virtual labs.
The reasons for higher demand for AI driven applications include the need for more virtual services, more expedient delivery of services, and filling the gap of staff shortages through the utilization of more automated workflows. Although all those reasons are quite valid, special attention should be given to the creation of artificial cultures while producing new AI applications and services.

How can our decisions when creating new AI applications affect the way in which societal power differentials, access to information, and selected metadata contribute to the creation of artificial cultures, of what we perceive to be valid or invalid, of labeling?
It has been stated that AI needs metadata and a more involved community of professionals addressing concerns like algorithmic bias, human interactions, and ethics in AI. What is our role as AI practitioners in creating artificial cultures? This session will provide a space for the exchange of ideas among concerned collaborators through the review of several case studies.

Priming Artificial Intelligence for Research in Information Sciences
Jessica K. Barfield (University of Tennessee- Knoxville)

This abstract discusses a methodological approach which can be used to prime artificial intelligence (AI) in studies which have applications to library and information sciences. Generally, priming occurs when an individual is exposed to a stimulus or experiment condition that may influence his or her subsequent response(s) in the study. As systems become equipped with more and more AI, information science researchers have developed an interest in determining how AI influences user behavior for a variety of tasks. However, for several reasons, such as the complexity, time, and cost of creating a fully functional AI system, researchers in information sciences may be reluctant to explore how the use of AI may affect user performance in their areas of interest. To work around these issues, in my research on evaluating user attitudes towards AI-enabled technology, I have recently begun exploring the methodology of priming to present participants with images that represent different levels of AI. In my view, priming is an example of a methodology that may create the impression that the technology of interest to an information science researcher is equipped with different levels of AI. Considering the method of priming, I propose a "realism scale" in the use of AI technology for information sciences research which from the low end consists of the participant being asked to think about (or visualize) an AI technology; next, the participant may be exposed to an actual image of the technology; then moving up the realism scale, the participant may be exposed to an animated video of AI-enabled technology; and lastly, but most realistically, the participant may directly interact with the AI-enabled technology. For the middle two methods, in my research using the priming method, participants are informed that they will be interacting with the AI technology for a particular task (which may require different levels of intelligence). By informing the user that the technology has different levels of AI, as I indicated above, we are priming the subject to react to the technology as if an AI-enabled entity. In my studies on the perception of AI enabled robots, I had subjects view the picture of a robot (with different appearances as an IV) and then had participants read a narrative (or hear the robot speak the narrative as another experiment condition) explaining its AI capabilities. In this research, priming occurred by nature of the capabilities the robot informed the user it possessed during the narrative (the AI techniques included natural language processing abilities, facial and emotion recognition abilities, and computer vision for navigation). In one study I performed on the evaluation of robot intelligence, a robot indicating that it possessed all the AI capabilities listed above represented the

high level of AI. Thus far, I have found that subjects evaluate a robot differently as a function of its perceived level of intelligence, and thus that priming is an effective methodology for user studies with AI technology when the research question can be explored without the use of a physical robot.

Predicting the Upcoming Topics and Actors of Fake News

Kevin Matthe Caramancion and Xiaojun (Jenny) Yuan (University at Albany, SUNY)

The role of AI in the battle against misinformation and disinformation has been integral. The distinction of disinformation from the former is the deliberate intent to deceive. The umbrella term fake news in social networking sites encompasses other forms of deception ranging from deepfake video renders to astroturf bots where AI typically functions as a robust enabler. On the other hand, AI has proven to be a key instrument for its management and control in the form of detection technologies ranging from visual-based to linguistic-based deceptions. The gap, however, lies in the fact that the applications of AI as a defense mechanism against deception are mostly saturated on mere detection rather than prevention in the first place. The cybersecurity principle of "think like a hacker" stipulates that threats are better controlled when precluded or anticipated, to begin with. In particular, AI technologies that are grounded on predictive prowess, rather than detection, warrant their own stream of research. The possible dimensions where deception prediction can be applied to include but are not limited to the (1) next probable subjects/topics of the upcoming fake news surge based on a current population discourse and, (2) the likelihood of users to be likely the actors of interest and concern in creating and spreading misinformation and disinformation content based on their social media behaviors. Although this promising domain may raise privacy and surveillance issues, its availability as an option should not be outright rejected. Finally, when used in conjunction with detection technologies, these preventative agencies may just be the key to a stronger wielding of AI against misinformation and disinformation.

Living and Working with Robots in the University Libraries

Aaron Choate, Elliott Hauser, and Katie Pierce Meyer (University of Texas at Austin)

"Robots in the libraries" will not be one thing, but many. How do we ensure that the adoption of robotics technologies across a range of library settings is supportive of libraries' missions, beneficial to patrons, and equitable for library staff? We propose to deeply research this question in a large academic library and articulate a replicable process of adapting robotic technology to specific user and administrative needs in library settings. The results of our work will inform the few but growing robotic deployments in libraries and, more importantly, position libraries at the forefront of the coming conversation around public interest robotics. We describe our evolving research program here to solicit feedback and potential collaboration opportunities with SIG AI members.

Step One Is Not: "Put Robots In The Libraries"

Too often technology projects race to deployment, before fully understanding the context within which deployment will occur. Our approach will involve 1-2 years of

ethnographic, historical, and participatory design work before robots ever enter the Libraries' workflow. We will aim to understand the people, institutional context, and opportunities stakeholders see for delivery robots within the complex sociotechnical fabric of library operations. Based on our team's professional experience and initial fieldwork, we believe that interbranch materials delivery is a promising use of autonomous service robots, but our methods are designed to surprise us by uncovering other opportunities. Additionally, a historical and ethnographic understanding of materials delivery as a practice, combined with our professional and research expertise in libraries, will allow us to propose policy changes alongside any eventual robotic deployment that will holistically maximize its benefits.

As we move from observation to participatory design, we will engage our robotics colleagues, experts in autonomous service robots, to collaboratively design, adapt, and deploy a pilot delivery robot. Our holistic methodology will continue in this phase, identifying any barriers to social or technical sustainability, and constantly revising the design of technology and socially conditioned practices that surround the deployment.

Perhaps most importantly, we will discover and articulate best practices for other institutions to replicate our process, rather than merely our results. This process-based approach could help guide the deployment of robotic technologies to achieve a range of goals in library contexts.

Centering Libraries in Autonomous Service Robotics Research

The PIs (Aaron Choate, Elliott Hauser, and Katie Pierce Meyer) are part of Living and Working with Robots, an eleven-PI project across seven disciplines, funded by Good Systems, a UT Austin Grand Challenge. As part of this team, we have discovered that our robotics colleagues view the long-term deployment of autonomous robots as a grand challenge of its own that demands deep collaboration. By combining robotics and social scientific methodologies, our team aims to produce a replicable method for deploying robotic research systems in specific contexts in ways that are ethical, helpful, and sensitive to the overlapping social contexts surrounding them. We're committed to positioning libraries as central to this inquiry, benefitting libraries as institutions while contributing to the emerging field of public interest technology more broadly.

Explainable AI and Information Provision

Martin Frické (University of Arizona)

Supervised Machine Learning (SML) is core to modern AI. What SML can tell us is that certain observable features, feature-data, are connected with other observable labels, label-data— for example, that being written by J.K. Rowling is correlated with being a popular book. Sometimes an SML will tell us the labels without being able to tell us of the features, or of the connections, that it used in the inference. In these cases, the SML is a black box. It is an Oracle. You give it a book, it will tell you if the book will be popular.

Libraries and information providers both run their own affairs and also act as intermediaries or gatekeepers between that which they provide and those they provide it to. The predictions of Oracle SMLs can be of different levels of seriousness in their consequences. Being wrong about whether a book will be popular would not usually bring into play death, destruction, or expense. But, in principle, an Oracle SML, could work on an entire country, every dietary item, every lifestyle, and every medical condition, and tell individuals what they should be eating and doing, perhaps even suggesting that some folk consume arsenic. Information providers need to be wary of Oracles.

Perceived Fairness of Facial Recognition System Deployment in a University Setting

Hengyi Fu (University of Alabama)

Facial recognition technology (FRT) is an emerging technology that will significantly transform our understandings and experiences of monitoring in a range of public and private spaces. While Facial recognition technology is now being introduced across various aspects of our life, the controversial nature of FRT and improper uses often generate critical concerns and even resistance. Current research on human interaction with FRT mainly focuses on individual level usage in private space; there is a limited number of studies focusing on users' perceptions of FRT usage in organizations such as workplace, they were mostly quantitative and did not capture the in-situ, nuanced human-surveillance technology interactions. In addition, limited research has been conducted specifically on organizational fairness issue associated with FRT.

To fill the research gap, our project aims to explore: What are individuals' lived experiences of facial recognition technology in a university setting? How organizational fairness is fulfilled or violated from users' perspective? To answer those research questions, we conducted a qualitative study in a public university in the United States, using semi-structed interviews. The study took place right after the implementation of a facial recognition system at a makerspace reconstructed from a library in a campus building. To our best knowledge, this is the first empirical study of how users perceive and interact with FRT in a university setting. Interviewees were asked about how they viewed and responded to the FRT implementation, what are some concerns they have, and how they viewed the organization that made this decision. In total eighteen participants were interviewed, inducing four faculty and fourteen students. Eight of them were affiliated with the makerspace, including one manager, two graduate assistants, and five undergrad interns.

Drawing from the lens of organizational fairness, we reported how participants constructed the perception of fairness and how the perceived fairness was violated by the university. Our findings show that the deployment of the facial recognition system in the makerspace has violated all three types of organizational fairness: distributive, procedural, and Interactional. In terms of distributive fairness, most participants felt that their privacy had been invaded, and their concerns of data security, possible bias, equal access, surveillance and possible abuse of authority were also not addressed or even being taken into account. They also felt the possible mismatch may negatively impact certain minority groups

disproportionately and hurt the community nature of the former library. From the procedural justice perspective, the decision-making process of implementing the facial recognition system failed to include diverse stakeholders. The administrators simply assumed that such a controversial technology would be eventually accepted and didn't even provide official training for student staff regarding the FRT system. Regarding interactional fairness, the administrators failed to show respect for the stakeholders and failed to provide sufficient explanations regarding how the system worked, why it should be installed, and how the decision was made. Our findings also suggested first impression counted, without pre-deployment user education, negative perceptions would remain and increase even for those who support the deployment.

The Affordances of AI for Everyday Information Practices
Noora Hirvonen (University of Oulu, Finland)

Artificial intelligence (AI) systems that collect, process, and act on data, have become embedded in the daily lives of people in digitalized societies. AI systems have been integrated into information intermediaries such as streaming and video-sharing services that use AI to recommend content; social media applications where AI is used to automatically filter messages and target advertisements; and search engines using AI tools to rank results. Moreover, AI technologies are used to power tools such as voice assistants, chatbots, and real-time navigation. People can also be subjected to AI-based profiling and automated decision-making. These AI-powered systems are already an integral part of the current information infrastructure, shaping the practices of seeking, evaluating, using, and creating information.

There is a growing area of human-centred research on AI generating crucial knowledge on making AI systems fair, accountable, transparent, and safe. However, this research tends to be centred on technical and legal concerns or specific platforms rather than the practices that are re-shaped as new technologies are being appropriated into them. A clear gap in research exists concerning the ways AI systems mediate everyday information practices, that is, the ways people seek, evaluate, share, use, and create information as part of their daily activities. Filling in this gap requires knowledge of both the technological features of these systems and of humans interacting with them in everyday situations.

One way to acknowledge how AI systems mediate information practices is to consider them as tools that, when appropriated by particular people in particular situations, come with different affordances. Affordances, when understood as relational rather than as properties of things, can be studied through the practices they become realized. Mundane and routinised everyday practices, in particular, can become invisible or largely unquestioned and unnoticed. Therefore, observational research methods can be particularly fruitful in unpacking the ways AI systems are enmeshed with and influence the everyday information practices of people in diverse situations, and with different needs and competencies.

As the amplified use of AI-powered systems is currently changing the dynamics in many spheres of human life, having an increasingly fundamental impact on people's actions and thinking, there is a clear need for furthering understanding of the impact and implications of these socio-technical developments. Including the perspective of information practices and exploring how AI technologies transform those practises is a research area where the information science field has an important contribution to make.

Affordances of AI for music-related information practices
Ville Jylhä (University of Oulu, Finland)

In recent years, artificial intelligence (AI) and algorithms have become a naturalized part of our everyday lives. Algorithms and AI shape our experience of the everyday as we delegate mundane routines, such as information searching and analysis, to be performed through algorithmic functions.

Like search engines and social media platforms, music streaming services also rely on algorithms and AI to power their software. In my PhD thesis I focus on music-related information practices and how algorithms and AI mediate those practices. People who have developed a serious leisure interest in music tend to be in search of new music and information about music they listen to, while amateur musicians also aim to share their own music online.

However, modern information mediators can also act as barriers and the users can only benefit from them if they adopt appropriate information practices. The affordances of cultural tools such as AI-powered technologies make certain actions possible, but at the same time, the affordances also shape the practices. Mostly due to the use of algorithmic selection and amplification of popular content in search engines and music streaming services, listeners face new challenges while trying to obtain new music suitable for their taste and musicians have hard time getting their own music heard. For example, the content people find online is often personalized by previous searches and to most users there is no way of telling how editorial playlists have been assembled and as the services widen their use of algorithmic engines, different forms of ambiguities, such as algorithmically produced tracks and ghost musicians, are likely to increase.

The data collection and analysis are guided by theoretical-methodological framework called mediated discourse analysis, which views social action and discourse as inextricably linked and seeks to work out a way to understand the relationships among human actions and discourses. Data collection will be carried out through interviews and participant observation to build up a picture of how the information practices are developed and influenced by affordances of the environment people operate within. Interviews and observation insights are compared to each other in order to point out possible contradictions and congruence between the participants' actions and talk.

My PhD research will contribute to filling in the gap in research concerning the affordances that AI applications have for amateur musicians' information practices. In order to develop deeper understanding of this relationship, the study seeks to pay attention to human needs and affordances for practices within a sociocultural environment, as Zhao, Zhang, Tang and Song (2021) have suggested. The concepts of mediation and affordances are used in the study to help explain how information practices are patterned by the technological, social and cultural construction of different artifacts.

References

Zhao, Y.C., Zhang, Y., Tang, J. and Song, S. (2021). "Affordances for information practices: theorizing engagement among people, technology, and sociocultural environments", Journal of Documentation, Vol. 77 No. 1, pp. 229-250. https://doi.org/10.1108/JD-05-2020-0078

Finnish young people's media literacy of deepfake

Yucong Lao (University of Oulu, Finland)

In the digital era, various technologies have caused difficulties for audiences to distinguish between the real and fake information in cyberspace. In the scope of inauthentic information, deepfake is a kind of burgeoning digital creation powered by artificial intelligence (AI), diffusing on vast social media platforms. More specifically, deepfake refers to a type of synthetic media in which a person's face in an image or video is replaced by someone else's, based on neural networks and facial mapping technology.

However, by using biometric information of human face, these AI-generated artifacts lead to more social issues involving disinformation and security, which might cause threats to our society. Meanwhile, video consumption has gradually been one of the most popular ways for young people's online entertainment, which indicates more deepfakes are possibly accessed by these digital natives. Therefore, for avoiding the potential risks, it is essential for us to attach importance on the young people's media literacy, including the thinking of deepfake and the competence of identifying deepfakes.

Rooted in a Finnish context, in my PhD research project, I will apply multi-method encompassing crowdsourcing and qualitative text analysis to explore Finnish young people's media literacy of deepfake from three perspectives: first, their consumption with deepfakes; second, their understanding and attitude towards deepfakes and awareness of hidden AI technology behind deepfakes; third, their competence of identifying deepfakes. In this manner, I can further figure out the power relations between deepfake and Finnish young people.

Through the investigation, this research can complement the research domain of AI and society from the aspect of deepfakes. As most of the previous studies about deepfakes focus on the technical side (e.g., the algorithms of deepfake and the deepfake detection techniques), this research is able to generate social scientific insights, from the side of audiences, contributing to new strategies of regulations targeted at deepfakes. Furthermore, the study indicates the new climate in the digital landscape and arouses people's attention on the authenticity of information.

Ethical Obligations in Big Data & ML Research Support
Casandra Laskowski (University of Arizona)

Libraries are increasingly providing data support, though the extent of the support varies from institution to institution. Services regularly include consultations and data management (like data repositories). While a great deal of conversation has revolved around the ethical concerns of ml projects and research (e.g., bias, dirty data, etc.), what obligation do libraries have in support of these efforts? There is a fundamental difference between hosting research papers and thesis in a repository and hosting data. Do we have an obligation to ensure data repositories have sufficient metadata and descriptions to inform future researchers of potential issues? Do we refuse to host suspected dirty datasets that might be used as a foundation for results that are harmful to vulnerable populations? Or do we feign neutrality and do nothing? Does the intent of the research behind the data matter? How do we draw the line? In consultations, are we obligated to inform researchers of potential pitfalls in their research methodology? If they ignore concerns, are we obligated to facilitate poor research?

One example that comes to mind is the use of historic crime data for hotspot analysis that informs the future allocation of police resources. Rashida Richardson, Jason M. Schultz, Kate Crawford noted that police data is problematic because it is built from "dirty policing," yet is still used for predictive policing efforts. If similarly problematic research is done at our institutions, do we have any ethical and moral obligations regarding consulting on such projects or housing such problematic data? And if we do, are we trained to fulfill them?

Using AI and ML to Optimize Information Discovery In Under-utilized, Holocaust-related Records
Richard Marciano (University of Maryland)

This talk describes the interdisciplinary work of archivists, scholars, and technologists to demonstrate computational treatments of digital cultural assets using Artificial Intelligence (AI) and Machine Learning (ML) techniques that can help unlock hard-to-reach archival content.

The project is an extended, iterative study applied to digitized and datafied WWII-era records housed at the FDR Presidential Library, rich content that is regrettably under-utilized by scholars examining American responses to the Holocaust.

The presentation details the benefits of interdisciplinary collaboration for evaluating user needs, identifying and applying tools and methodologies (including ML through object detection and AI through Named Entity Recognition or NER) , and reaching the real-world outcome of public access to augmented data.

A Meta Framing Problem How Can A.I. Conceptualize Search as Learning?
Alamir Novin (University of British Columbia)

A major obstacle for A.I. research is referred to as the "framing problem" (Chow, 2013; J. Fodor, 2006; J. A. Fodor, 1987; Minsky, 1997; Shanahan, 1997). In brief, the problem asks how A.I. can differentiate relevant information from an environment to achieve common sense? For example, if an A.I detects a student's cellphone ring in a classroom, followed by the student apologizing to the teacher, but then the teacher has a sudden heart attack – how can the A.I. frame the first two events as correlated without correlating the third? This paper begins by asking before finding the correct frame, can A.I learn from the search for the correct framing? Can an A.I. learn from the journey before the destination?

Simple framing problems are addressed by teaching A.I. to learn instructions (e.g., non-monotonic formalisms (Shanahan, 1997)). With more common frames, A.I. needs to learn how to learn a frame, such as when A.I. attempts to emulate a learner's plasticity (Kovalev & Yakasova, 2020; Milano et al., 2020) using a cascade correlation model. However, more sophisticated frames wrestle with questions of epistemology and ontology (Chow, 2013; Dennett, 1984; J. Fodor, 2006; J. A. Fodor, 1987; Tannen, 1993). For example, how can an AI system think about how it thinks (i.e., metacognition)? Metacognition is what enables humans to not just complete a search task but conceptualize the search as a learning experience (i.e., Search as Learning (SAL) (Cole, 2020; Rieh et al., 2016)). This paper experiments on the framing problem and concludes that just as cognition requires framing, metacognition requires a metaframe (i.e., where an agent can metacognate about frames). Conversely, without recognizing a metaframe, the framing-biases in an environment are adopted and never challenged. As this paper's experiment found, if people lose track of the meta-frame they exhibit cognitive biases. The paper concludes that SAL is affected by how the metaframe is recognized – presenting a problem for both A.I. and human learning.

Experiment

Two groups of students (N=50) were asked to use a mock search engine that retrieved identical results for each group. The top of the search page framed cellphones on campus as an educational issue while the bottom framed cellphones as a health issue. Both groups were asked to search cellphone risks for university students. However, the second group's task was framed with the phrase 'The health of students is important to the university.' Of interest was how the groups recognized the frames.

Results

Three sources of framing-bias are found in each frame: Source-bias emerged in the environment frame due to how different sources framed cellphone problems, algorithmic-bias emerged in the medium frame due to students trusting the search page algorithm, and cognitive-bias affected the subject's frame due to heuristics (e.g., whether to explore the breadth of a search page first before clicking on a link or to explore the depth of the first few results at the top first led to different understandings). The first group's metaframe perceived 'cellphones on campus as a classroom distraction while the second group's metaframe perceived them as a student mental-health issue.

Analysis

The two groups' metaframes were influenced by whether students recognized the embedded frames. Most importantly, within both groups, some participants mitigated the framing-biases. These latter participants mitigated framing effects by recognizing the influence of the three frames in their meta-frame. In conclusion, the brain is not 'in a vat' and depends on the framing of others as much as one's own frame. It is imperative for agents to both distinguish and make value judgements comparing the influence of the three embedded frames in a metaframe.

The metaframing problem applies to A.I. systems. A.I. must recognize all three frames to mitigate biases. Take for example an A.I robot searching for a frame with minimal confirmation bias: First, the robot must search whether its neural network confirms a prior frame (e.g., a bias towards a predetermined reward). However, even if the robot has no prior biases, it must search that the medium by which it chooses to interact with an environment is framed to confirm biases. Finally, the robot must also search that its data sources are framed by external agents in agreement with a bias (e.g., a cultural bias). Finding a more accurate frame in a framing problem requires an agent to learn from its search using a metaframe. The real challenge is if an A.I can meta-frame the complex problems that humans struggle with during SAL. Thus, the metaframing problem becomes a problem of framing other agents' framing problems.

The Impact of AI on Information Technology Workforce
Sang Hoo Oh (Florida State University)

Artificial Intelligence (AI) is considered as one of the most important and influential technology in the coming decades. It is expected to have an impact on every aspect of our society and economy. AI is also expected to transform the nature of existing jobs and have a huge impact on the future workforce in various industries. The impact on jobs includes both job creation and displacement, and increased labor productivity to widening skills gap. Consequently, AI has recently fueled interest in future of work debate as the technology achieved superhuman performance across a wide range of economically valuable tasks. This research explores the impact of AI on future information technology workforce. The information technology workforce had not been affected by previous technologies. However, AI is different

from previous technologies that it can be applied to highly educated and well-paid jobs. Our review of related research suggests that little is known about the extent to which AI will have impact on the information technology workforce. Hence, this research proposes the new model to examine the impact of AI on the competency skills and knowledge of the information technology workforce. In this model, we apply natural language processing algorithm to measure the overlap between verb-noun pairs of the U.S. Department of Labor's job competency model and verb-noun pairs of AI patents. AI patents contain key application and essential information about AI technology, and the U.S. Department of Labor's job competency model describes the knowledge and skills needed at different levels of information technology jobs. The outcome of this study will provide empirical evidence to the extent to which information technology workforce's knowledge and skills would be impacted by AI. It will also provide insights to the information technology workforce on how they should prepare for the era of AI.

A Case Study Using NLP to Analyze LibChat Transcripts
Meng Qu (Miami University)

Natural Language Processing (NLP) is one of the fields in machine learning that has the ability of a computer to understand, analyze, manipulate, and potentially generate human language. It is widely used for machine translation, text simplification, sentiment analysis, and many language processing situations. With the power of NLP, the researcher was able to analyze and extract useful information from digitized word-based data, such as online library service chat transcripts.

LibChat is the built-in chat reference module for LibAnswers from Springshare. It is widely used for online chat services in libraries. In regular online service hours, anyone who reaches the online chatbox can ask the librarian on duty. Although LibAnswers provides a comprehensive statistics dashboard, telling the quick facts of metadata statistics summary, it cannot fetch critical information from transcripts. In this case study, the researcher wanted to solve a few questions: what are the top frequently asked questions in their initial questions and among the conversations, and the typical conversation depth/turns have made in one service; what are the typical patrons' groups or affiliations; and if the frequency of online services is affected by school events (i.e., the beginning of semesters, mid-term exams, final-term exam, vacations, or so). Besides, the researcher wanted to generate a data visualization report to indicate the findings.

With these questions, the researcher applied a series of machine learning programming for transcripts analysis in Jupyter Notebooks with Python3. The main processes include: download and clean the data, to extract only the related information; applying Natural Language Toolkit (NLTK) and TextBlob for information retrieval (text tokenization, simplification, tagging, classification, etc.,); using Python Natural Language Processing Library (PyNLPl) to complete the tasks of n-gram extraction and modeling; applying Polyglot for sentiment analysis; finally, using Pandas and Seaborn and a set of related algorithms to visualize the findings into graphic charts. Based on the raw data of nearly 7000 records during the past

two years, some interesting findings came with the analysis. Moreover, the outcome has testified that online chat services are related to school events.

This case study is an experiment of applying NLP-related technologies for online chat transcripts analysis. In the library field, we can apply the findings in multiple aspects, such as building an auto-answering chatbot, improving virtual reference services management, and facilitating library instruction topics to address the frequently asked questions. In the future, there should be more data collected, such as in-person reference conversations, to help expand the research outcome.

Folk Theories and Explainable AI (XAI)
Michael Ridley (Western University)

Can folk theories about how AI systems work enhance explainable AI (XAI)? Folk theories (aka mental models) are the beliefs users have about how things work. These beliefs and heuristics shape their expectations, inform their actions, and influence their trust. They "need not be technically accurate (and usually are not), but they must be functional" (Norman 1983).

XAI is a set of strategies, techniques, and processes that include testable and unambiguous proofs, various verification and validation methods that assess influence and veracity, and authorizations that define requirements or mandate auditing. While the audience for an explanation can be system developers (primarily interested in performance), clients (primarily interested in effectiveness or efficacy), domain experts (primarily interested in work related outcomes), or regulators (primarily interested in policy implications), it is the needs of everyday users (primarily interested in trust or accountability) that are the most urgent.

By eliciting the folk theories of users of AI systems it is possible to enhance XAI by aligning explanatory capabilities with user perceptions.

Addressing Bias and Fairness in Search and Recommender Systems
Chirag Shah (University of Washington)

Bias is omnipresent -- from data to algorithms, and from framing of a problem to interpreting its solution. In this talk, I will highlight how such bias in general with machine learning techniques, and in particular with search and recommender systems cause material problems for users, businesses, and society at large. The examples span areas of search, education, and health. I will then introduce the idea of marketplace as a way to find a balance or fairness in the system and address the issue of bias, among other things. I will draw specific examples from our work on search and recommendation systems to demonstrate that achieving fairness in a marketplace and addressing bias in data and algorithms are not just morally and ethically right things to do, but could also lead to a more sustainable growth for various industries, governments, and our scientific advancement.

Information Science Informing Ethical AI

Ali Shiri & Toni Samek (University of Alberta, Edmonton, Alberta, Canada)

In 2020, the International Federation of Library Associations and Institutions Governing Board agreed upon the September 2020 IFLA Statement on Libraries and Artificial Intelligence. In December 2020, the Association for Information, Science and Technology, alongside the Association for Library and Information Science Education and the iSchools, released a Statement on AI ethics and the contributions of diverse voices in the discussion. These statements reflect a social responsibility to advance rigorous scholarship and practice in the face of the prolific contributions to a topic that has captured the attention of the global academic enterprise, the information professions, and intercultural society.

A recent search on the Scopus database on AI and Ethics shows a 700% increase in the number of publications in this area, from 330 in 2018-2019 to 2184 on August 26, 2021. Importantly, in relation to the increasingly multidisciplinary nature of the discourse of AI and Ethics, is the finding only 24.4% of the publications appear in computer science literature and with over 24% in the social sciences and humanities and several other disciplines. In a topic and trend study of AI and Ethics, the number of published documents peaked at 330 in 2018 with data and information among the key concepts in the discussion of AI and ethics. The data suggests the scholarship on AI ethics is closely linked to the established fields of data ethics, computer ethics and information ethics. Many of the data and information ethics−related themes and topics (privacy, confidentiality, trust, and moral principles and ethical concerns) have frequently appeared in the literature of AI.

The rapid growth of literature on AI and ethics is significant, revealing scholarship reflective of emerging, complex, sensitive, and multifaceted implications for research and development and educational purposes, including the information science curriculum and new approaches to information science teaching and learning. In this presentation, we make two key arguments: First, data and information are the foundational constructs in any discussion of ethical AI. Library and Information Science has a long history of research into information organization, analysis and processing; information retrieval; information interaction, cognition, and search behaviour; library and information ethics; and the fundamental concept of relevance. The domain of Library and Information Science has developed a solid knowledgebase for understanding data and information and how humans interact with them. It is interesting to note how the current literature on ethical AI is surfacing the importance of what library and information science has been researching for over 50 years, namely the ethical treatment of information and data and how data and information is created, processed, understood, retrieved, interacted with, and acted upon. Secondly, we argue informed discussions of AI and Ethics should draw upon both the well-established areas of information ethics and library ethics, and closely examine and conceptualize the key ethical principles that have been developed over several decades. This will benefit constructive (not just productive) contributions to the development of a framework for understanding, teaching, and learning of data-focused and information-focused ethical AI.

Reluctant and Non-Library Users: Is a chatbot the answer?
Tienya Smith (Queens Public Library)

Public library chatbots like Calgary's Scout engages readers and helps them to find library materials. But what about reluctant readers or non-library users? Would an AI solution like a chatbot or intelligent assistant motivate them to read or use the library?

Relevancy and discoverability are essential in capturing the attention of non-library users. Recommenders or intelligent assistants that could help reluctant users find timely materials or specific library services could potentially attract these individuals to the library. Additionally, embedding chatbots or intelligent assistants on the library's social media channels could make these AI solutions more discoverable to this hard-to-reach population.

During my lightning talk, I will closely examine these solutions and their relevance in an urban market that serves disconnected youth, reentry adults, and low-income families.

AI Fairness and Algorithmic bias
Emmanuel Sebastian Udoh (University at Albany) and Abebe Rorissa (University of Tennessee- Knoxville)

Recent advances in artificial intelligence (AI) and machine learning (ML) have triggered a rising proliferation of automated decision systems, leveraging the data generating power of the Internet of Things (IoT) and the power of predictive analytics.

Particularly, predictive analytics enables not just the discernment of patterns in large datasets, but also predictions, prognostications, and forecasts. Expectedly, AI/ML technologies are increasingly being deployed across various domains including healthcare, advertising, stock market investments, among others. However, in the last decade and a half, AI/ML analytics are increasingly being deployed in public sector domains to make and automate decisions with far reaching consequences for persons and groups, such as in parole decisions, credit/loan eligibility, healthcare coverage eligibility, recidivism decisions, college admissions, teachers' performance evaluations, public housing eligibility, and even predictive policing. While the appeal of AI/ML technologies lies in their touted affordances of efficiency, accuracy, and the elimination of human bias (Burrell, p.5), but while efficiency and accuracy might have anecdotal evidentiary support, the elimination of human bias has remained a pipe dream. The new technologies are largely perceived to be exacerbating the situation, as they tend to digitalize or codify human bias (digital redlining), raising a plethora of concerns and challenges, including fairness, inclusion, privacy, accountability, transparency and equity, among others. In fact, the new technologies have been labelled 'instruments of oppression' of minority demographics, people of color, disabled persons, women, and even older adults (Eubanks, 2016; O'Neil, 2016; Noble, 2018).

So far, the focus of scholarship has been on conceptualizing or demystifying the underlying algorithmic opacity (Sandvig, 2015; Oswald et al., 2018), the need for more responsible practice through accountability and transparency from AI/ML vendors and practitioners, and the urgency for more experimental studies. However, the threat of 'algocracy' (Danaher, 2016; Zambonelli et al., 2018; Steiner, 2012) portends a multidimensional problem including the legal, contextual, political, social, technological, policy, among others. While no technology is completely free of human bias (Gillespie, 2017; Crawford, 2016), they can at least be fair, and fairness has become more and more nuanced (Pessack & Shmueli, 2020; Kearns and Roth, 2020; Barocas & Selbst, 2016) and too often unaddressed in discourses on algorithms, AI and ML. Accordingly, this presentation aims to underline the importance of a reasonable expectation of fairness by users, especially the vulnerable and often marginalized, and consistent with organizational values, societal laws and regulations (Accenture Federal Services, 2018). We will provide an opportunity where scholars (and practitioners) in AI/ML can dialogue on the key challenges of the new technologies from a human standpoint; such issues include fairness, inclusion, the different kinds of bias potentially doting the algorithmic decision systems landscape, the role of data quality as a potential game changer (Barocas & Selbst, 2016; Vagle, 2016), and the proclivity for pernicious feedback loops or self-fulfilling prophecies. The broader impacts of these challenges, novel concerns and challenges will also be discussed, all aimed at centring fairness in the AI/ML discourse and recommending new paths towards a fairer AI/ML.

References

Accenture Federal Services (2018). Responsible AI: A Framework for Building Trust in your AI Solutions (accenture.com), https://www.accenture.com/_acnmedia/PDF-92/Accenture-AFS-Responsible-AI.pdf. Accessed: 2021-05-30.

Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. Calif. L. Rev. 104 (2016), 671.

Burrell, J. (2016). How the Machine 'Thinks:' Understanding Opacity in Machine Learning Algorithms. Big Data & Society January–June 2016: 1–12.

Crawford, K. (2016). Can an Algorithm be Agonistic? Ten Scenes from Life in Calculated Publics. Science, Technology & Human Values, 41(1), 77-92.

Danaher, J. (2016). The Threat of Algocracy: Reality, Resistance and Accommodation. Philosophy and Technology Vol 29 issue 3, pp. 245-268.

Eubanks, V. (2016). Automating Inequality: How High-Tech Tools profile, Police, and Punish the Poor. New York: St. Martin's Press.

Gillespie, T. (2017). Algorithmically recognizable: Santorum's Google problem, and Google's Santorum problem. Information, Communication & Society, 20(1). http://www.tandfonline.com/doi/full/10.1080/1369118X.2016.1199721#abstract

Kearns, M., Roth, A. (2020). The Ethical Algorithm: The Science of Socially Aware Algorithm Design. Oxford: Oxford University Press. Pp.91-93.

Noble, S. (2018). Algorithms of Oppression: Race, Gender and Power in the Digital Age. New York: NYU Press.

Noble, S. (2012). Missed Connections: What Search Engines Say about Women. Bitch magazine, 12(4): 37-41.

O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. New York: Crown.

Oswald, M., Grace, J., Urwin, S., & Barnes, G. C. (2018). Algorithmic risk assessment policing models: lessons from the Durham HART model and 'Experimental' proportionality. Information & Communications Technology Law 27(2): 223-250.

Pessach, D., & Shmueli, E. (2020). Algorithmic Fairness. AEA Papers and Proceedings, 108, 22–27.

Sandvig, C. (2015). Seeing the Sort: The Aesthetic and Industrial Defense of 'The Algorithm.' Journal of the New Media Caucus http://median.newmediacaucus.org/art-infrastructures-information/seeing-the-sort-the-aesthetic-and-industrial-defense-of-the-algorithm/

Steiner, C. (2012). Automate this: How algorithms took over our markets, our jobs, and the world. New York, NY: Portfolio.

Vagle, J. L. (2016). Tightening the Ooda Loop: Police Militarization, Race, and Algorithmic Surveillance, 22 MICH. J. RACE & L. 101, 102.

Zambonelli, F., Salim, F., Loke, S. W., De Meuter, W., & Kanhere, S. (2018). Algorithmic Governance in Smart Cities: The Conundrum and the Potential of Pervasive Computing Solutions. IEEE Technology and Society Magazine, vol 37, issue 2, pp. 80-87.

It takes a village: Building Mason-Library's Orientation Conversational Agent through conversational marketing and human machine interaction techniques using a multi-agent approach

Trevor Watkins (George Mason University)

Orientation events present one of the greatest opportunities for academic libraries to connect with new students and faculty to market library resources and services. The pandemic forced all departments at George Mason University, including Mason library, to transition orientation sessions from in person to online. In person sessions While assessment of orientation sessions is usually in the form of soliciting feedback after the session is complete, it is not a great indicator of how much

information is retained. During a recent analysis of some of our virtual chat sessions, we discovered that many questions posed by users could have been addressed during orientation. While most libraries today use chatbot technology to optimize research and information services, the goal of this conversational agent is to reduce orientation-related inquires, so that virtual reference can take on more research and reference focused inquiries.

In this presentation we describe the process of putting together a team of librarians, communication and marketing specialists, software engineers, and graphic and usability designers and briefly discuss their responsibilities. We introduce the creation and implementation of a chatbot technology readiness level measurement system (based on NASA's TER measurement system) that guides the project workflow. Finally, we discuss how we are using conversational marketing and human machine interaction techniques that incorporate librarian and intelligent agents to build a multi-agent conversational agent.

Designing Responsible AI Systems for Older Adults: Opportunities and Challenges

Xiaojun Yuan, Bahareh Ansari, Mehdi Barati, Benjamin Yankson, Kevin Caramancion, George Berg, DeeDee Bennett Gayle (University at Albany, SUNY)

According to the Census Bureau, all baby boomers will be 65yrs or older by 2030 (Census, 2019). This will increase the demand for Artificial Intengece (AI) aided systems, which can handle repetitive tasks, to complement and relieve the workload of health professionals and caregivers. Such demand calls into question the use, access, and protection of sensitive patient information requiring the urgent need for research that addresses issues about privacy and other ethical considerations.. Older adults are usually underrepresented in usability studies of new technologies and therefore, their special needs and preferences were often not met by new technologies (Paez & Del Río, 2019; Aguirre & Abadía, 2017). The Federal Trade Commission (FTC) reports show that older adults are an especially vulnerable population to online fraud and identity theft (Federal Trade Commission, 2020). In particular for AI technologies, older adults can be vulnerable to nudges to over share their personal information and compromising their information security. Responsible AI is a framework ensuring the ethical, transparent and accountable use of AI technologies in a way "consistent with user expectations, organizational values and societal laws and regulations" (Accenture Federal Services, 2018). Responsible AI concerns the ethical and legal perspective of AI, which become even more important when AI serves for older adults. For example, how can we design an automated system (e.g., the social robot) that may reliably assist older adults whilst respecting their emotions, privacy, culture, beliefs, and dignity? This research aims at raising awareness of the importance of this topic, and exploring possibilities of connecting researchers from various fields together to contribute to AI governance to ensure AI systems can develop with privacy and ethics in mind to serve older adults and satisfy their information needs. In particular, we will explore the possible challenges and address them. Instead of designing algorithms or prototypes for AI systems, we focus on how older adults use, and interact with such systems, and their concerns and perceptions of using such systems, as well as

design implications that we can provide for system designers and developers. Possible future research collaborations will also be discussed.

Keywords: Older Adults, Ethics, Privacy, Technology Use, Aging, Artificial Intelligence, Responsible AI

Census (2019). By 2030, All Baby Boomers Will Be Age 65 or Older: 2019. https://www.census.gov/library/stories/2019/12/by-2030-all-baby-boomers-will-be-age-65-or-older.html#:~:text=Since%20then%2C%20about%2010%2C000%20a,of%20the%20U.S.%20Census%20Bureau. Accessed: 2021-05-30

Accenture Federal Services (2018). Responsible AI: A Framework for Building Trust in your AI Solutions (accenture.com), https://www.accenture.com/_acnmedia/PDF-92/Accenture-AFS-Responsible-AI.pdf. Accessed: 2021-05-30

Paez, L. E., & Del Río, C. Z. (2019, July). Elderly Users and Their Main Challenges Usability with Mobile Applications: A Systematic Review. In International Conference on Human-Computer Interaction (pp. 423-438). Springer, Cham.

Aguirre, D. F., & Abadía, I. (2017). Review of accessibility and usability guidelines for website design for the elderly people. Sistemas y Telemática, 15(42), 9-29.

Federal Trade Commission (2020). Consumer Sentinel Network Data Book for 2020. Washington, DC: Federal Trade Commission. Available: https://www.ftc.gov/system/files/documents/reports/consumer-sentinel-network-data-book-2020/csn_annual_data_book_2020.pdf