NSF Undergraduate Education (DUE): grant 1941782

#STEMaiCMU

Carnegie Mellon University
Mechanical Engineering

ARTIFICIAL INTELLIGENCE
and the Future of STEM and Societies

Workshop Report

Prepared by Conrad Tucker, Ph.D.

conradt@andrew.cmu.edu

Table of Contents

Workshop Overview

The *digital age* has the potential to transform the way that society teaches, learns, and prepares for the 21st century career landscape [1]. The workshop on *Artificial Intelligence (AI) and the Future of STEM and Societies* (AI+STEM Workshop) brought together a diverse group of experts in AI and STEM, and non-experts, to discuss ways that AI could enhance the way that STEM education is taught in academic institutions, industry, and government. The diversity of people and topics represented at the workshop is a testament to the broad interest in AI, STEM, and workforce development across different sectors and disciplines.

Workshop Objectives

On December 2nd and 3rd, 2019, 146 participants across academia, industry, and government convened on the campus of Carnegie Mellon University (CMU) to:

i) discuss ways that artificial intelligence (AI) could enhance the efficiency of Science, Technology, Engineering, and Math (STEM) knowledge acquisition and workforce retraining
ii) identify industries and tasks that are the most at risk of AI automation and discuss ways to mitigate the negative perceptions of AI that may develop as a result
iii) discuss ways that the fields of AI and STEM education could be made more diverse in an effort to mitigate the biases that may result in AI integration in STEM education and workforce development
iv) establish a national network of stakeholders passionate about AI and STEM education so as to inform policy makers and society in general about the current and future trends in AI and STEM

Workshop Format

The National Science Foundation (NSF)-funded workshop attracted a diverse group of participants from numerous disciplines including education, engineering, computer science, psychology, philosophy, among others. The two-day workshop's format provided keynote speeches on each day of the workshop, several panel sessions focusing on specific aspects of AI and STEM education, and breakout sessions, during which workshop participants could each contribute to the outcomes of the workshop by participating in specific workshop themes centered on AI, STEM and society.

Figure 1: (left) CMU Provost James Garrett Welcomes Workshop Participants at the Sunday GMSP Evening Celebration Event; (right) GMS Scholars in STEM fields across academia, industry and government alongside GMS Program Officials

A pre-workshop event was held on Sunday, December 1st, 2019, that recognized the achievements of the Gates Millennium Scholars (GMS) Program in advancing the pipeline of STEM talent across the U.S. and the globe, with a particular focus on students from underrepresented backgrounds. The $1.6 Billion grant from the Bill and Melinda Gates Foundation has resulted in over 20,000 GMS alumni during the course of the 20-year program [2]. Workshop participants were welcomed by CMU Provost James Garrett and GMS STEM Scholars (Figure 1), highlighting CMU's commitment at the highest levels, to the aims and objectives of the workshop and its participants. The workshop sought to overcome two critical misconceptions: (i) that AI exists solely within the computer science domain, and (ii) that STEM is a term that relates more to K-12 education. The reality is that AI, both its development and application, extends far beyond the computer science domain [3], and that STEM education impacts learners of all ages and domains, including adult learners who may be interested in reskilling and career development [4].

Participating Organizations and Demographics Data

Table 1 outlines the organizations represented at the workshop. Figure 2 presents the distribution of professions represented at the workshop. As can be seen from Table 1, academia had the highest representation, followed by industry organizations, non-profits/non-governmental organizations, and federal/government organizations.

Table 1: Workshop Representation Across Industry, Academia, Government

Academia	Industry	Non-Profit/Non-Government Organization	Federal/Government
Air University	Candence	Allegheny Health Network	Air Force Office of Scientific Research (AFRL)
Air Force Institute of Technology	Covestro	Atlantic Council	Army Research Laboratory (ARL)
Carnegie Mellon University	CNH Industrial	Bill and Melinda Gates Foundation	Office of Naval Research (ONR)
Carnegie Mellon University-Africa	Dun and Bradstreet	Carnegie Bosch Institute	National Aeronautics and Space Administration (NASA)
Carnegie Mellon-King Mongkut's Institute of Technology Ladkrabang (CMKL) University-Thailand	Energid	National Academy of Engineering	National Institutes of Health (NIH)
Florida Agricultural and Mechanical University	Google	RAND Corporation	National Science Foundation (NSF)
Northwestern University	IBM	Susan G. Komen Foundation	U.S. Army Futures Command
Oregon State University	Intel	The Brookins Institution	
Penn State University	Lockheed Martin	United Negro College Fund (UNCF)	
State University of New York at Buffalo	Microsoft	United Nations International Telecommunication Union (ITU)	
University of Arkansas	NVIDIA		
Wright State University			

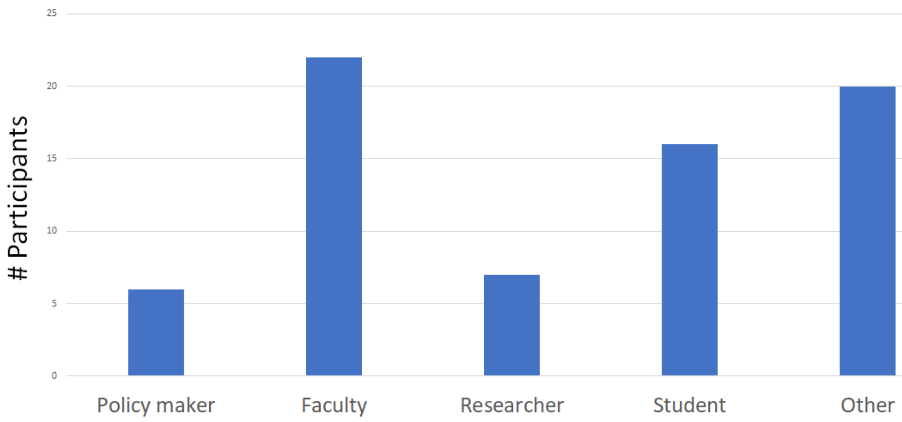

Figure 2: Distribution of Professions Represented at the Workshop

While representing a range of professions, participants tended to be 1) relatively new to their job role, with approximately 42% reporting 0-5 years in their job role, or 2) highly experienced, with 36% reporting 16+ years in their job role. A pre-survey question asked workshop participants to define AI in their own words. This question was included on the survey because definitions can reveal insights into how someone might, for example, feel unsure about the role AI can play in STEM and society. It is interesting to note how these perspectives, from different work sectors and with differing job longevity, resulted in similar key understandings when defining AI. Participants tended to include these five keywords in their definitions: “AI,” “human,” “computers,” “data,” and “tasks.” Figure 3 provides a distribution of these key words. Most notably, the word human (18) appeared as frequently as AI (18). This may help guide researchers to consider the human dimension of AI at the very early stages of AI theory and development.

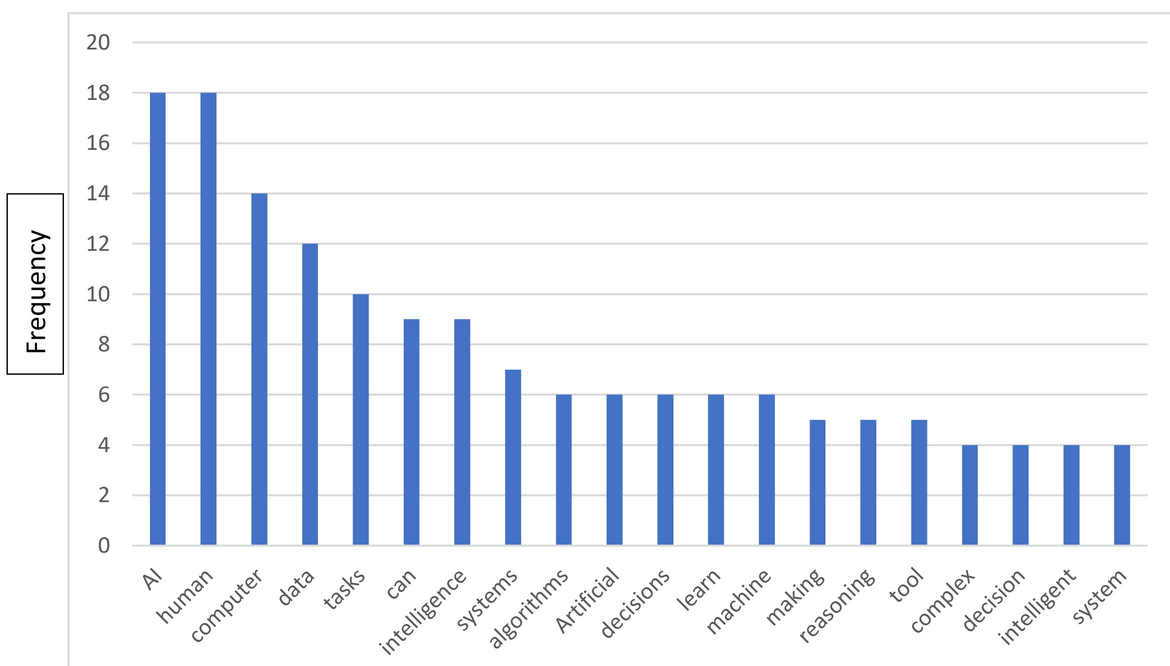

Figure 3: Distribution of Key Words Used to Define AI

While the responses in Figures 3 and 4 indicate that participants don't all operate with the same definition of AI, their foundation is generally the same. If we use their key words to define AI, it appears they see AI as an extension of humans and not just a machine. If we see AI as a collaboration between people and technology, does this perspective help us to recognize the potential in AI and to not fear it as something that will adversely take over our lives? The visualization in Figure 4 offers a Wordle representation [5] of participants' definitional words revealing the overarching theme of intelligence and the connecting of humans and computers.

Figure 4: Word Cloud Visualization of Participants' Definition of AI

Key Workshop Themes and Discussion Topics

AI and STEM and the Need for Academia, Industry, and Government to Work Together

Key workshop themes were addressed during keynote talks, panel sessions, and breakout sessions. Figure 5 represents the conceptual link presented at the start of the workshop to demonstrate the interconnectedness of academia, industry and government, and how that interconnectedness impacts the advancement of AI, STEM, and societies. The mechanical gears in Figure 5 served as a metaphor for this interconnectedness, which also represents the stereotype that is typically used to convey the mechanical engineering discipline [6]. Unlike actual gears in mechanical systems, it was quickly pointed out during the workshop that the metaphor in Figure 5 realistically looks less like a functioning system of gears that symbolize industry, academia, and government, and more like isolated gears that have little interaction with one another. *"The system is broken,"* expressed a

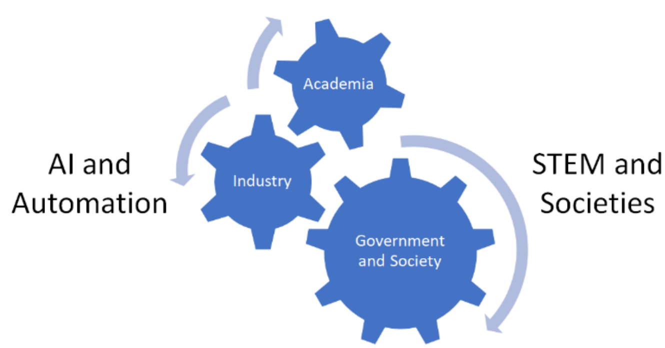

Figure 5: Mechanical Gears Metaphor Introduced at the Start of the Workshop

panelist during one of the panel sessions, a sentiment that received significant applause from participants in the audience. Participants' perceptions of isolated gears spinning in Figure 5 may help explain why such a sentiment may exist across academia, industry, and government. These domains are each successful at what they do, but workshop participants found that there was little collaboration among them that translated to real-world impact and outcomes. Participants noted that in academia, success is typically measured by churning out graduates that then enter the workforce as industry or government employees, or as next generation academics who continue the cycle. In industry, companies hire talent produced by academia in order to maximize shareholder value through the sale of products and/or services. Government also competes with industry for talented individuals who will help advance the safety and security of the nation. Some participants felt that these entities, or gears, often feel disconnected from each other, constantly optimizing their own objective functions, with values that may be misaligned.

Some workshop participants indicated that tenure-track faculty in academia are focused on particular incentives: attaining tenure, the next research grant, or academic accolade. Students' metrics are even more limited and are often reduced to a single value—the grade point average (GPA). The focus on student GPA as a measure of success was dispelled through several industry studies presented at the workshop. These internal studies, while company-specific, found little to no correlation between an individual's GPA and their career success. Furthermore, some industry experts present at the workshop have even explored the correlation between where an individual attained their undergraduate degree and their subsequent success in their profession. It may come as a surprise to some, but the same speaker revealed at the workshop that there was also a weak correlation in this dimension within their organization.

Workshop participants expressed concern while discussing the industry perspective on metrics that maximizing shareholder values can, in some cases, also lead to short-term thinking and decision making that negatively impacts the health and safety of society. Workshop participants noted that our government is also plagued with misaligned incentives, as evidenced by the quarterly jobs report as a measure of how well our government is performing, which in many cases, does not capture a comprehensive, longer-term vision, especially in communities that are less represented in terms of their voice on the national stage [7]. Some participants pointed to the use of misaligned metrics as a contributing factor to some of the socio-economic challenges faced in society, expressing concern about such metrics being integrated into AI systems that may cause these challenges to scale. For example, some workshop participants were surprised to learn of the disparities that exist in the valuation of homes across racial lines [8]. Research from one of the workshop panelists reported on the billions of dollars in lost value of homes owned in African American neighborhoods, even when factors such as crime, social status, etc., were controlled. Yet, as one speaker brought to our attention, inclusion and diversity are tied to a successful society.

It is against this backdrop: metrics, competing objectives, and potentially misaligned values, that workshop participants were called upon to set everything aside, and focus on the future of AI, our STEM workforce, and how our nation and the world at large, will evolve in the coming years. The unifying factor at the workshop was that no domain, be it academia, industry, or government, was immune to these potential disruptions. Consequently, workshop participants were highly motivated to freely share ideas and outline a set of action items and guidelines pertaining to AI and STEM.

AI for Skill Automation, Not Job Replacement

One workshop discussion theme centered on the concern that many of society's perceptions about AI include "Terminator style" machines that are either here to destroy humans or "take our jobs." The process of job destruction and job creation is nothing new and existed long before the inception of modern-day AI [9]. During the workshop, participants discussed today's fears about AI and automation replacing jobs. It was recognized that more must be done to better demystify AI to the general public. One idea resonated with the participants: more local community outreach events to connect researchers in academia and industry with individuals who may/may not be well-versed in AI. The robotics community represented at the workshop is well aware of the challenges of deploying robotic systems alongside human operators in, for example, manufacturing facilities (Figure 6). On one hand, human operators are needed to demonstrate the positive impact that automation has in their workspace. On the other hand, there are fears expressed by human operators that demonstrating such positive impact may result in the inevitable replacement of their job by automation. Some workshop participants expressed concern about the potential of creating an atmosphere of distrust. While some industry experts at the workshop proposed focusing on task automation—rather than job replacement—resolution could not be reached in terms of how task automation, while not necessarily resulting in complete job replacement, would not impact wages. If one takes the overall wage earned for a job as a summation of individual tasks directly relating to that job, and then a subset of those tasks being automated by AI, automation has the potential to deflate wages or, at the very least, keep them stagnant. While there were healthy discussions around the topic of AI, automation, and its impact in shaping the workforce, pressing questions remain about how to maintain high wages in the advent of skill automation and how to rapidly retrain workers whose skillsets are obsolete.

Figure 6: Humans Working Alongside Robots

Advancing STEM Education and Workforce Reskilling through Emotional-AI

This workshop theme explored ways that AI could replicate human emotions to motivate and better engage students during the learning process. We can look to the multibillion gaming industry that is extremely successful in creating an emotional connection between humans and pixels on a screen with which humans can engage with for hours on end [10]. The success of the gaming industry has even spurred a relatively new dimension of research called "gamification," wherein elements that make a game engaging (i.e., badges, points, leader boards), have been adapted to other domains (healthcare, education, etc.) to motivate individuals towards higher task performance or knowledge acquisition [11]. Some workshop participants stated that advancing the domain of emotionally engaging AI was critical to engaging students in any meaningful, prolonged manner. Meanwhile, other participants were more skeptical and did not agree that humans should create AI that either replicated human emotions or resulted in humans being emotionally attached to AI. Participants expressed concern than an emotional connection may be difficult for learners to separate from, especially if this emotional connection is customized down to the individual level. Furthermore, participants expressed ethical concerns regarding how early to introduce such an AI in the student development process. Would emotional connections with AI compete with the emotional human-to-human connection? This domain of research, while potentially impactful, raises some safety and ethical concerns regarding exploitation and misuse. I.e., if AI is able to understand what motivates one to learn and make certain decisions, should you the human, know how it has learned to do that? In other words, should the AI be opaque and focus on the overall objective (help a learner achieve new knowledge or skills), or should it be transparent in how it is able to achieve this? Do these goals necessarily have to be competing objectives, or can society advance towards both explainable and functional educational AI? These and similar questions were proposed as some of the next steps in building on the foundations of the workshop and the stakeholders involved.

AI and its Potential Impact in Shaping Muscle Memory, a Key Ingredient for Building a Sustained STEM Workforce

How long does it take to master a skill? Consider that same skill and ask how long it takes for that skill to degrade to levels that are dangerous—both to the individual performing the task—or the individuals who may be impacted by the task being performed. Questions like these posed during the workshop resulted in many fruitful discussions. Muscle memory, defined by Merriam-Webster as “the ability to repeat a specific muscular movement with improved efficiency and accuracy that is acquired through practice and repetition,” was a key topic of discussion at the workshop [12]. Beyond teaching students new STEM knowledge, participants expressed interest in exploring how AI could potentially help the workforce retain the skills and knowledge that they had already acquired. Participants responded to these questions by stating that these questions depend on the precision needed for a given skill and the consequences that result in said skill degrading. One of the workshop speakers explained that the “learn-by-doing” component of building muscle memory plays a critical role in a wide range of high-performance tasks currently needed in the workforce, ranging from fighter pilots—whose muscle memory may start to degrade within hours/days of performing a repetitive task—to an AI programmer, who may be able to resume training an AI agent even after months of taking a break. Participants asserted that novel technologies that combine AI and simulation must be integrated into the workforce training efforts to maintain muscle memory at levels that enable optimum human performance. This need by industry and government begs the question in terms of how academia can help advance basic science into what causes the degradation of muscle memory and whether that process can be slowed down or even halted? The implications of this scientific breakthrough would have direct impact on society, potentially saving billions of dollars spent each year on retraining and recertifying the workforce. An outcome of these workshop discussions was the realization that more basic research was needed into the science of how to retain the knowledge/skills gained by learners, rather than having to constantly retrain them.

Protecting AI from Biased and Corrupt Data Sets

Given the data-driven approach to many of today’s state-of-the-art AI algorithms, the workshop spent a substantial amount of time exploring solutions to ensure the veracity of data. Workshop participants shared personal stories about existing human biases exhibited in housing prices, access to healthcare, aviation, and the quality of education. Participants expressed concern that if human biases are not resolved in the critical systems that govern society today (i.e., education, healthcare, transportation, etc.), potential exists for AI systems designed by humans to also be prone to biases. A recent United Nation’s study discussed during the workshop, found that digital assistants, which dominate our modern digital society, may be perpetuating gender biases by

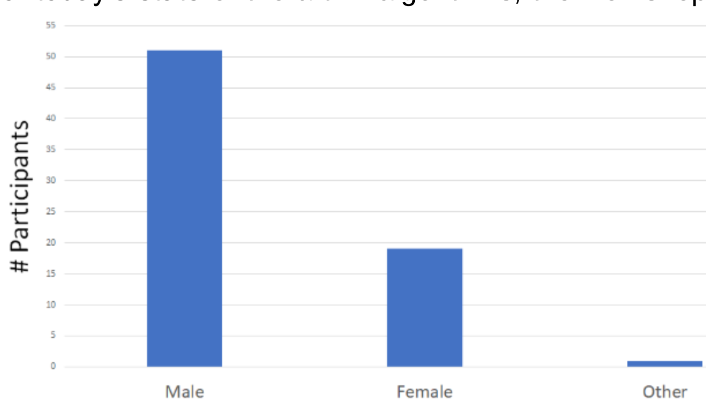

Figure 7: Gender Distribution of Workshop Survey Respondents

attaching gender and qualities such as obedience, to digital platforms [13]. Regarding gender biases, the workshop distribution itself demonstrates the challenges of attaining representative participation across different demographics. In Figure 7, of the 71 workshop participants who responded to the gender question, 51 were male, 19 were female, and 1 "Other." The process of inviting participants included reaching out to leaders and decision makers within organizations (e.g., deans/department heads in academia, division leaders from industry and government) and asking them to recommend individuals whom they thought would benefit from participating in the workshop. In addition to unequal gender distributions, it was noted during the workshop that there exists a significant disproportion of AI research being advanced by groups that are not demographically representative of society. These differences can already be seen in facial recognition systems that have been shown to have more difficulty recognizing darker faces than lighter faces [14]. The implications of these biases in STEM and the workforce are profound. For example, for online learners taking an examination that uses facial recognition for identification, darker-skin students may be put at a disadvantage if the system is unable to recognize them, or assigns an attribute to them that may further advance biases. Figure 8 presents the race distribution of workshop survey respondents, which is not necessarily reflective of the current distribution found within AI research and development teams. Participants recognized the need to ensure representation exists across a wide range of demographics to minimize biases, an ongoing issue that needs to be addressed across all sectors.

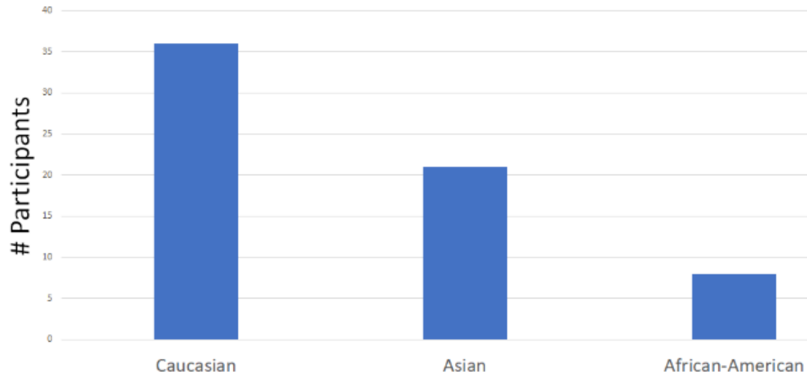

Figure 8: Race/Ethnicity Distribution of Workshop Survey Respondents

Perhaps more of a threat to STEM-based AI systems are threats not readily observable. Facial recognition's inability to detect a dark skin student during an exam may cause significant issues, but it represents a bias that can be readily observable. However, the advent of generative neural networks has enabled hyper-realistic data to be generated that is indistinguishable from real data [15]. From a STEM perspective, the risks associated here are the potential for students to seek out publicly available content (e.g., a video demonstrating how to use a widget) that has been manipulated by AI systems. AI-driven tutor systems that are trained using publicly available data are also at risk of being corrupted by malicious data. Furthermore, this manipulation of data may not be detected until it has been widely disseminated (e.g., integrated into a STEM curriculum that has been viewed by thousands of students).

Workshop participants discussed efforts by policy makers to hold for-profit-organizations, such as social media networks, accountable for the spread of misinformation. Experts hold mixed opinions as to whether for-profit organizations should—and even could—address the problem of misinformation. It was noted by workshop participants that this shift in policy may have the unintended consequence of transforming for-profit companies into the guardians of free speech. Furthermore, since the scale and scope of misinformation detection far surpasses the human bandwidth available in today's society, it was highlighted that organizations are relying on AI algorithms to serve as the filter for misinformation detection. As an optimistic look into the future, workshop participants expressed confidence in the democratic system of western societies in having their populations participate in the discussions about the future of AI and data accessibility, as opposed to a top-down approach. Participants were optimistic that while less efficient than more hierarchical systems, a democratic process would result in more sustainable outcomes.

Workshop Breakout Sessions

In addition to the general workshop discussions that took place during the course of the two-day workshop, targeted breakout sessions were organized on each day to provide participants with a more focused STEM/AI topic to discuss, with a smaller subset of participants. Across a total of six workshop breakout sessions, participants were randomly partitioned and assigned to one of each of these breakout sessions. In order to attain a cross pollination of ideas, participants switched their breakout session assignments during each day of the breakout sessions. This also helped keep participants engaged by not repeating the same breakout session for each day.

Breakout Session 1: AI for Advancing Personalized Learning

This breakout session explored the potential benefits of the National Academy of Engineering (NAE) Grand Challenge of *advancing personalized learning* [16]. The breakout session discussions began with several participants drawing the distinction between personalized content delivery—and the actual process of learning being personalized. Participants noted that AI systems are already being used to personalize the delivery of learning content through existing, adaptive AI systems. This personalization has been made possible via recent advancements in AI chatbot systems and intelligent tutoring systems, in addition to access to high quality educational content freely available on sites such as YouTube and Coursera. Participants argued that personalized learning, however, was much more difficult to achieve, compared to personalized content delivery, since it involves the assessment of whether an individual learner has met the learning objectives and has acquired the desired concept/content. Participants highlighted several key challenges, such as identifying and addressing a student's learning needs and establishing helpful two-way exchanges between students and teachers/existing AI educational systems. It was noted that proper metrics currently do not exist for evaluating how well students learn concepts and apply conceptual thinking. While metrics such as quizzes and exams are the default standard in academia, several participants believed AI-driven personalized learning could also expand the diversity of assessing student learning outcomes.

Several participants believe a balance in AI education is essential, between self-discovery of knowledge and prescribed pathways, such as lectures that are teacher-directed. Some participants suggested AI can be a tool to engage learners and provide them more personalized learning. For example, learners can be divided into groups that choose specific projects, instructionally guided and supported with AI. It was also emphasized that achieving diversity of teams needed to be central to an AI learning system, while maintaining a foundation of common knowledge among the group members. When talking of passive learning (e.g. AI in a classroom), the importance of ensuring common background knowledge and experience across learners was highlighted. Some participants believe that the paradigm of physical classroom instruction was outdated, and that AI has the potential to transform the way students learn through digital media. While there was great enthusiasm of AI and the digital transformation of learning, several participants expressed skepticism in the ability of AI to completely make the human instructor obsolete. Several participants pointed to the impact that coaches and mentors have in motivating students to learn. On the issue of motivation, participants highlighted the fact that there currently exists an abundance of books and other educational learning content, many of which are free to access and that do not necessarily translate to the acquisition of knowledge. While AI may be able to personalize content for students, some participants were skeptical of AI's ability to motivate students to want to learn. That, some participants felt, was where humans could play a critical role. In this paradigm, AI would be responsible for advancing the generation of content that could help personalize learning, while humans could serve as motivators when students either become stuck when learning content or become less motivated to learn at different stages of the learning process. There was disagreement among participants as to whether this human engagement, had to be in person, or whether this engagement could also be digital through technologies such as video conferencing/virtual reality.

The breakout session on personalized learning outlined tangible action items moving forward that could help facilitate the advancement of AI's role in achieving the NAE Grand Challenge of *Advancing Personalized Learning*. Participants agreed that no matter what form AI takes in helping to advance personalized learning, the data available for AI to train upon, has to be constantly evolving and current. Teachers' knowledge of AI systems also needs to be state-of-the-art, for them to be aware of the strengths and limitations of AI systems. Participants agreed on the idea that learning as a whole has to expand beyond the existing focus on k-16, and instead, include the "life-long learning" paradigm, wherein an AI agent could be seen as a digital tool that evolves with the individual learner. It was noted that other countries have already deployed adaptive learning systems (some of which were developed here at CMU) to assist in the increased demand for educational content. As a result, the US should survey the evolving landscape of educational AI systems in order to remain leaders not only in the development of AI, but also in its deployment and utilization. It was noted that US regulatory policies need to be revisited to ensure that they facilitate the deployment of AI in educational settings, while vigilantly maintaining student privacy and data security needs.

Breakout Session 2: The 21st Century "Classroom"

What does the 21st century "classroom" look like? The answer may not be a physical classroom at all. "Students do not need to come to school," "things can be done online remotely," and "there is no need for a physical classroom," were some of the sentiments expressed by participants during this breakout session. However, not all participants believed that physical human involvement was unnecessary in the teaching and learning of content; some participants argued that teaching students how to learn was more important than what students learned. Participants also noted the importance of the generational disparity between seekers of knowledge, who may be younger and more digitally connected, and providers of knowledge, who may be older and more comfortable with traditional, in-class approaches to teaching. Nevertheless, society is experiencing a transformation of the learning paradigm, with a growth in the abundance of freely available, high-quality content on sites such as YouTube and other digital media. Participants discussed what areas within the current educational system would be most impacted by this disruption. Participants agreed that AI has the potential to advance personalized learning, such as how and when to help students, based on each learner's capabilities. Evidence of adaptive learning systems, such as on-demand video content recommendation systems, already exists. However, in order to be applicable to STEM and learning, advancements are needed in terms of the data types and learning contexts provided to students. Unfortunately, existing video and text platforms for STEM content delivery lack the level of personalization that would make AI adaptive to each learner. Rather than seeing AI as a substitution for human instructors, some participants believed that the human-human interactions were necessary for advancing knowledge acquisition. As a result, AI could be perceived as aiding the human learner to understand each student (e.g., using an iris detector to predict students' engagement in a course topic) and recommending time/context interventions that human instructors could provide to each student. This was a point of debate among participants, focused on how much personal interaction with an instructor was needed for learning. For AI-based education to scale, many argued that society is going to have to accept a reduction in the availability of human-human learning.

As society moves towards a more digitally connected paradigm, participants questioned whether learning would shift from more individual-based learning and assessment to more team-based assessment. Some participants expressed concern that if AI is focused on one-on-one instruction, it may result in creating "knowledge bubbles," wherein learners indeed learn STEM content, but lack the knowledge on how to apply that STEM knowledge when dealing with other humans, or exposed to diverse situations. Regarding diversity, some participants argued for the inclusion of "arts" in the advancement of technology in the 21st century (i.e., STEAM, instead of just STEM). The argument for STEAM is that the human creativity dimension of learning is going to be humanity's key differentiator when it pertains to co-existing with AI and automation, and that STEM should not be advanced in the absence of arts and creativity. Some participants went onto argue that the focus on STEM leaves out other critical domains, such as economics and philosophy, that AI must be able to understand and communicate, as many scientifically advancements begin with posing philosophical questions, and then shift towards becoming more scientific over time with experimentation and data.

Action items discussed included the need for stronger partnerships between industry and academia, in an effort to connect learners with more real-world experiences, similar to the apprenticeship paradigm. As AI becomes more prevalent, participants recommended that math skills across society be strengthened so that AI is seen not as "magic," but instead, as a series of mathematical operations applied to data. Participants recommended that society engage more with the humanities fields, as they will play a critical role in how AI impacts the 21st century.

Breakout Session 3: Broadening Access to STEM through AI

This breakout session focused on how AI could help broaden access to STEM in ways that humans are unable to because of a wide range of challenges: geography, scarcity of teachers relative to learners, costs, etc. AI bots are a prevailing technology that participants felt could help expand access to STEM, ranging from children to adult learners. For children, conversational AI bots could broaden access to STEM knowledge by responding to early curiosity-types of questions in an interactive way (e.g., how are the colors of the rainbow formed?). For adult learners, AI bots that serve as the STEM analogous to Quora (a crowdsourced, Q&A platform), could help expand the STEM knowledge available to learners. Furthermore, it was discussed that AI bots may help reduce some of the trepidation that some students face when asking questions relating to STEM, while not trying to appear unintelligent to their peers and/or teacher. Enabling AI that allows students to "make mistakes" was deemed a critical component of what was needed to broaden access and engagement in STEM. There was great optimism expressed by participants in AI's ability to personalize the learning experience so that more people could be introduced to STEM content early in their lives in ways that related to their individual experiences. Participants felt that this, in turn, could potentially mitigate biases that currently exist in how STEM is currently taught (e.g., using a car or a bridge example as the default use case for teaching a wide range of STEM-related concepts, despite the fact that not all learners may connect with those examples). Along the personalization thread of the discussions, there was interest in reforming the evaluation metric of *grades* that are currently used to evaluate students' current or future capabilities in STEM. Rather than grades, it was proposed that AI has the potential to discover students' strengths in ways that elucidate their potential strengths in STEM, rather than grades, which are easier for educators to assess in a scalable and efficient manner. Participants felt that this change could help broaden access.

Participants believed that robotics could help foster interest in STEM and teach STEM concepts. In essence, this involves promoting student learning and interactions by taking the intelligence of AI bots and merging it with hardware. This approach highlights a potential scalability issue for physical AI. While STEM AI bots may be developed and disseminated to anyone with an internet connection, achieving such scale with physical AI in the STEM domain, may present many of the challenges that students from underrepresented groups face today. I.e., the digital divide due primarily to costs and accessibility. Broadening access means reaching low income students who may not have access or the means to interact with AI in ways that may lead them into STEM-related fields. Despite these potential challenges, there was significant interest among participants in having AI that could be integrated into things that students use and bond with, such as a doll. John Seely Brown and Paul Duguid's *Social Life of Information* was referenced during discussions to remind everyone of the importance of the closed knit relationship between technology and society, as well as how those relationships serve as a critical factor in the exchange of information.

Given that these automated systems may eventually serve as the guide toward or away from STEM fields, the participants concluded this breakout session by exploring biases in society to ensure that these biases do not translate to STEM AI systems. Starting with the interface in which learners interact with AI systems, researchers and designers need to think broadly in an effort to promote diversity and inclusion. Diversity in research and designer teams may help ensure that issues pertaining to biases are part of the AI STEM development process, and not necessarily addressed after an issue involving AI and STEM discrimination has already occurred. Participants discussed potential ways to overcome biases in AI STEM systems, such as the deployment of AI tools across a wide range of different learning environments, having role models play an active role in the development of STEM AI, and having job descriptions focus more on features that make an individual successful throughout the course of their career, not necessarily only specific to a job, as the job landscape is ever evolving.

Breakout Session 4: Data Ownership in the Age of AI

Data is defined by Merriam-Webster as "information output by a sensing device or organ that includes both useful and irrelevant or redundant information and must be processed to be meaningful" [17]. This breakout session discussed what it meant for humans to increasingly serve as the system being sensed to capture meaningful data. Several questions are motivated by the term *meaningful*, as it implies value to an entity or stakeholder. Trust of the data collector, and the data itself, became a central discussion topic during this breakout session, as participants considered the assessment of the value of data and whether knowledge for the social good, trumped individual preferences. Participants discussed whether individuals should be obligated to share their data if the aggregation of data from individuals results in better societal outcomes, such as a more efficient way to teach STEM. While individuals already relinquish certain individual preferences (e.g., obeying the speed limit) for a greater societal good (e.g., safer highways), some participants expressed concern in the lack of transparency regarding the societal benefits that stem from the sharing of their personal data. Participants even questioned the feasibility of having a *data property* paradigm, wherein individuals would be responsible for managing their own data across different entities and objectives. One of the challenges pertaining to data and its ownership that participants expressed was the question of what entity would be responsible for the acquisition, storage, and mining of a person's data. For example, the motivations for for-profit organizations may differ from the motivations for government organizations in terms of the utilization of individuals' data. A wide range of successful enterprises have been built around the mass collection of individuals' data, which is then sold to advertisers for more customized advertising. Some participants argued that since users accept the terms of service for certain platforms, they themselves bear the responsibility for the current state of the policies surrounding individuals' ownership of personal data. For example, if a user signs up for a "free" webmail account or social media platform, they should also consider the value proposition of the entity offering this "free" service and why such an entity would be motivated to do so. Regarding freedoms, both for organizations to collect and utilize data, and for users and groups to reject the collection and use of data, several cultural differences exist. During the breakout discussions, there were cultural differences expressed in terms of different regions in the U.S. and their trust in different organizations. In the U.S., some participants felt that more trust seems to be placed in for-profit organizations and their mass collection and utilization of user data, compared to the government attempting to do the same. Participants contrasted that to Europe, where for-profit organizations and their collection and utilization of user data has been met with increased skepticism and scrutiny. Some workshop participants expressed the opinion that it was too little, too late in terms of trying to protect users' data. Part of the challenge stems from the advancements of AI algorithms themselves, as new capabilities continue to be developed that discover insights from data. For example, an innocuous selfie video, while taken to teach a STEM topic, could now reveal certain health related aspects (e.g., pulse rate estimation from simply a video recording [18]) about the individual themselves, without need for additional data collection.

The breakout session concluded with an acknowledgement that discussions about data ownership needed to be a top national priority and communicated to the general public in a more palatable manner so that well-informed decisions could be made. There is an acknowledgement that the general public is still not well-informed of the value of their data. From a STEM education perspective, there was recognition that there exists tremendous potential in leveraging existing education and career trajectory data to assist individuals in choosing their fields of study or career paths. Participants expressed commitment to collaborating on future initiatives in this area and recommended that the set of upcoming opportunities be set as targets for collaboration.

Breakout Session 5: Ethics of AI in STEM and Society

Participants in this breakout session met to discuss the ethical challenges of wide-scale AI deployment in STEM and society and propose potential solutions that may mitigate these ethical challenges. While AI has the potential to enhance the efficiency of STEM knowledge acquisition, participants were quick to highlight concerns that AI may actually amplify already existing biases in society, which may result in further inequalities. This concern arises from the current paradigm wherein AI algorithms use training data that is generated directly from humans (e.g., via Social Media platforms) or sensors that were created by humans (e.g., image data acquired from camera data on a phone). Participants mentioned that if AI uses data generated by humans, data ownership becomes complex. Furthermore, there exists varying policies around the world pertaining to data ownership such as the *right to be forgotten* policy [19] that enables individuals living in certain regions in the world to remove their personal data from the public domain if they feel that it may be harmful to their image. Several ethical considerations stem from such policies, such as whether it is ethical to degrade the performance of AI (by individuals removing their data from the public domain) that is meant to teach students about a historical event that may paint certain individuals in a negative light.

In order for AI to advance beyond training from data generated by humans, participants explored the possibility of AI learning about the world through experiences in the physical environment, similar to how humans learn. Participants questioned whether it is ethical to even have a class of AI that mimicked the many dimensions of human intelligence, which includes emotions such as love, anger, etc. During one of the panel sessions, a panelist from industry made the case for AI that developed emotional connections with humans, as a way to foster the desire of humans to learn. Participants discussed whether creating emotional connections with AI should be a two-way street i.e., should only the human form these emotional connections with an AI, or should society create AI that can also form emotional connections with humans?

The two-way communication between humans and AI would inevitably mean that human data was being constantly transferred to a system that the AI further trained upon to refine its connection with its human counterpart. Participants questioned the ethical implications of this constant transfer of human data to an AI data processing system. These ethical concerns are even more critical in STEM fields, wherein students engaging with an AI system may be minors. Furthermore, if AI uses this data to refine its algorithms for more engaging interactions with students, at what point does this data collection process become human subjects research? Currently, many organizations involved in human subject research have an Institutional Review Board (IRB) that reviews new projects (or modifications to existing projects) to ascertain the ethical implications of the project. The IRB review process can take weeks, and in some cases, months to complete. With a dynamic AI that adds/removes/enhances features of its architecture, will the engagement with a student be placed on hold each time a modification of the AI features occurs in order for an IRB review process to be performed?

Participants concluded the session with a set of action items, including the creation of a white paper primarily focused on the ethics of AI in STEM and society. Participants outlined three layers of ethics that need to be explored: i) direct effects, ii) vulnerabilities, and iii) societal. For direct effects, these include situations wherein AI actions result in accidents that cause harm. Harm could range from administering an incorrect exam to a student, due to an incorrect facial recognition classification, to the loss of life, due to AI recommending an incorrect procedure for a student learning to fly a plane. The ethical considerations for AI vulnerabilities include the deliberate hacking of an AI system that results in the harmful examples above. From a societal perspective on ethics, participants questioned whether it was ethical for AI to replace human jobs, resulting in mass unemployment. Participants would like there to be more discussions at the national level that include policy makers, who would be responsible for enacting laws that ensure that AI is used in ethical ways that benefit not only a subset of society, but society as a whole.

Breakout Session 6: The Future of Work and Lifelong Learning

This breakout session explored the relationship of the future of work, education and learning, and AI technologies. It addressed three questions: How may AI technologies displace workers by taking jobs? How may AI technologies be used as educational tools to enhance human learning capabilities? And, how may the educational sector respond to the changes needed in education and in learning technologies? The AI technologies considered were the current generation based upon statistical learning also known as the 2nd wave and future AI technologies, which have the capability of “understanding” context or explainable AI known as the 3rd wave.

Participants indicated that there is a need to take a step back to reflect upon the nature of work in the context of AI technologies such as the clarification of “knowledge work” and “physical work.” They reflected upon several questions: What constitutes a job, and how may AI technologies change it? Will future AI technologies increase the efficiency of people learning new knowledge and skills that will enable them to perform their current job more effectively and become better educated to learn a new job?

One concern is how AI and automation would displace human workers through the substitution of AI technologies. The concern about machines displacing workers is not new and has been a societal issue since the beginnings of the Industrial Revolution, such as the Luddite movement. Machinery substituted for labor, leading to gains in productivity because of more output per worker, ultimately led to the creation of more jobs through economic growth with an overall increase in societal wealth. Now we must question whether AI technologies are different from industrial or non-AI computer technologies because of its impact on employment growth, and whether there will be a net increase in jobs. Also, it is unclear to which degree AI technologies will increase the effectiveness of life-long learning, such that new knowledge and skills may be acquired and developed so that people may learn how to perform new jobs.

Participants discussed the idea that jobs are composed of tasks. The group focused its attention on how the performance of tasks would change with advances in AI technologies, as well as the role of work to human purpose and meaning. Participants agreed that the threat of job loss not only exist at lower wage jobs, such as stocking supermarket shelves, but that AI technologies have the potential to replace jobs that are more intellectually demanding, such as teaching mathematics or designing scientific experiments. While a job in its entirety may not be eliminated, there would likely be changes in the number of workers needed—fewer workers for a specific job. And, there may be changes in the wages or return to labor with a fraction of the return increasingly accruing to an AI technology owner. Consequently, there are income distributional effects. Participants discussed how a focus on tasks, rather than jobs, may enable individuals to better prepare to acquire new skills for performing new tasks, which would redefine the work and job. It is unclear how workers, either knowledge or physical workers, will interact with AI technologies for retraining in knowledge acquisition and skill development. One possibility is the widespread availability of cost-effective AI-enhanced simulator technologies.

Participants expressed concern that organizations are ill-equipped in quantifying or pricing the value for each task. E.g., how much is one line of code worth to an organization versus a keynote talk? Who gets to decide? Clarification of the value of tasks and transparency of how different people get compensated, for example the marginal differences in compensation between a chief engineer versus a CEO, could generate organizational and social tension, but it could also lead to more equitable distribution of compensation across organizations and society.

While AI technologies have demonstrated significant capabilities for tasks, such as voice and image recognition, that can substitute or complement human work, for example decreasing the need for telephone customer service employees or complementing the work of physicians in the visual detection of tumors in CAT scans, participants thought that critical thinking abilities of humans, which is part of the quest for general artificial intelligence, would remain a human task for the foreseeable future. For certain tasks that are well-structured and that do require human judgment, such as playing chess or Go, AI technologies have proven superior. It remains unclear if this machine capability implies use in less well-structured circumstances, as does the question of how humans will team with AI technologies.

The breakout session concluded with a set of action items and a book recommendation, Andrew Grove's "*Only the Paranoid Survive: How to Exploit the Crisis Points That Challenge Every Company."* Grove's six forces framework illustrates the different forces that affect a business. One of the forces is the "possibility that what your business is doing can be done in a different way," which is a force relevant to AI technologies. Participants thought that more emphasis is needed on "lifelong learning," and that a "paradigm shift" is needed, one that recognizes people will require continuing education throughout their careers. It is a societal moral imperative to promote and enable the education of the whole workforce. For organizations, either public or private, a well-educated work force is a means to greater productivity and employee satisfaction and advancement. The economics of AI technology and how the technology will disrupt the business models of educational institutions remains an ongoing issue.

Workshop Statistical Analyses and Take-Aways

Participants' Expectations of the AI+STEM Workshop

Prior to the workshop, participants answered survey questions (on a 1-5 Likert scale) related to their expectations of the AI +STEM workshop. The stem question was "r*ate the extent to which you agree or disagree about the following expectations you have for the workshop*."

We conducted a cluster analysis to characterize their expectations of the workshop. Four expectation questions were: "I would like to build relationships and network with others," "I would like to gain more knowledge about AI in STEM education," "I would like to learn more about funding opportunities focused on AI in STEM education," and "I would like to learn more about research opportunities focused on AI in STEM education."

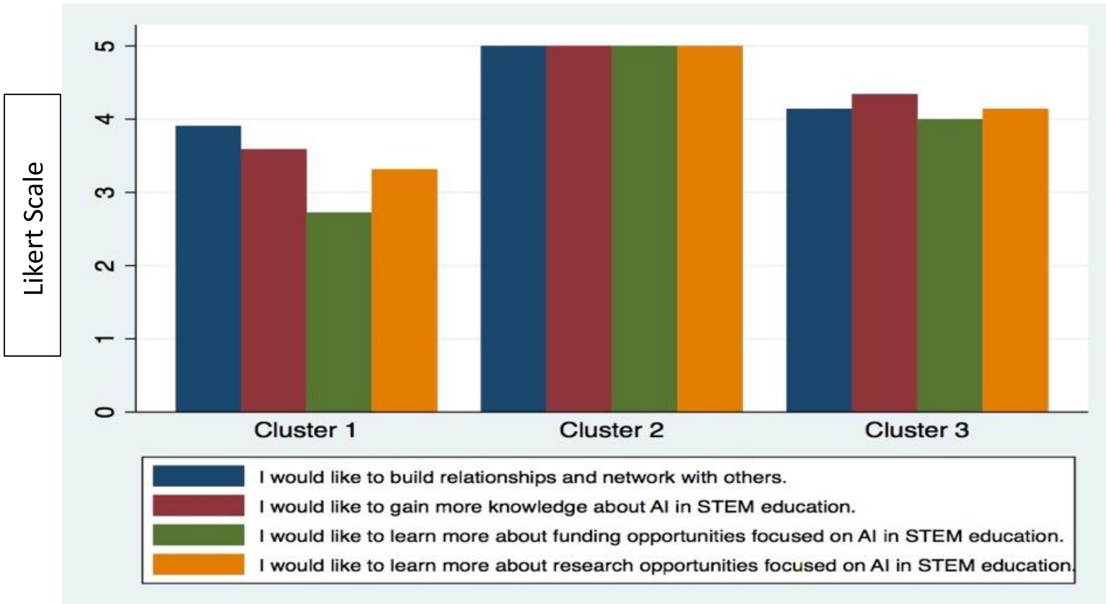

Figure 9: Cluster analysis on participants' expectations of the AI workshop

The cluster analysis revealed (see Figure 9) that the participants' expectation of the AI workshop can be categorized into three groups/clusters. We assigned descriptive, if unwieldy names, to these clusters by assigning high (mean >= 4.5), middle (4.5 > mean >= 3.5) and low (mean < 3.5) designations to *networking, knowledge, funding opportunity, and research opportunity*. Table 2 shows the percentage of current job setting (e.g., government, higher education, industry, other) in each cluster.

Cluster 1 is labeled "MidNET/ MidKnow/ LowFUND/LowREA" to reflect strong expectation for networking (mean = 3.91 on a 5-point scale) and knowledge (mean = 3.59) over funding opportunity (mean= 2.73) and research opportunity (mean = 3.32). Cluster 2 is labeled "High for All" to all 5 on a 5-point scale for networking, knowledge, funding opportunity, and research opportunity. This cluster is the smallest of the four. Cluster 3 is labeled " Middle for All " and reflect expectation for networking (mean = 4.14), knowledge (mean = 4.34), funding opportunity (mean= 4.00), and research opportunity (mean = 4.14).

Table 2: Chi-Square Test of Independence (Cluster and Current Job Setting)

		Current Job Setting				Total
		Gov.	Higher Edu.	Industry	Other	
Cluster	1	40.0%	19.6%	50.0%	100.0%	32.4%
	2	40.0%	17.4%	8.3%	0%	16.2%
	3	20.0%	63.0%	41.7%	0%	51.5%
Total		100.0%	100.0%	100.0%	100.0%	100.0%

Note: Sample size of cluster 1,2,3, and total are 22, 11, 35, and 68 respectively. Sample size of Current Job Setting Gov., Higher Edu., Industry, and Other are 5, 46, 12, 5, and 68 respectively.

Participants' current job setting in Government is overrepresented in Cluster 2 (High for All) and underrepresented in Cluster 3 (Middle for All). Participants' current job setting in higher education is underrepresented in Cluster 1 (MidNET/ MidKnow/ LowFUND/LowREA). Participants' current job setting in industry is overrepresented in Cluster 1 (MidNET/ MidKnow/ LowFUND/LowREA) and underrepresented in Cluster 2 (High for All). Other current job setting participants are overrepresented in Cluster 1 (MidNET/ MidKnow/ LowFUND/LowREA) and underrepresented in Cluster 2 (High for All) and 3 (Middle for All)

According to chi-squared test of independence, participants from the government sector had more individuals with higher expectations in all aspects (e.g., networking, knowledge, funding opportunity, and research opportunity) of the AI+STEM workshop. The majority of participants have higher education jobs (46 out of 68) and more than 80% of the higher education participants fall into Cluster 2 (Higher for All) and Cluster 3 (Middle for All). Meanwhile, industry participants fall into a profile of low expectations in funding and research opportunities, but middle level expectations in networking and knowledge. In other words, higher education participants have higher expectations of the AI+STEM workshop. There was no significant association between other participants' characteristics (e.g., gender, ethnicity, year of experience) and the clusters.

Participants' Satisfaction with the AI +STEM Workshop

After the AI+STEM workshop, we asked participants to answer survey prompts related to their satisfaction of the AI+STEM workshop. The stem prompt was "*rate your level of satisfaction with each of the following aspects of the AI workshop*."

We conducted a cluster analysis to characterize their satisfaction with the AI+STEM workshop. Five satisfaction questions were: "Through the workshop, I gained a deeper understanding of AI and its future impacts," "My expectations entering the workshop were met," "Information I learned about AI has improved my understanding of AI," "My networking at the workshop with others in the emerging field has led to opportunities for continuing conversations," and "I am ready to commit to becoming part of a sustainable network of workshop attendees who address AI issues." We labelled each criterion as the following terms: Future Impact, met expectation, Understanding AI, Networking, Participating AI community.

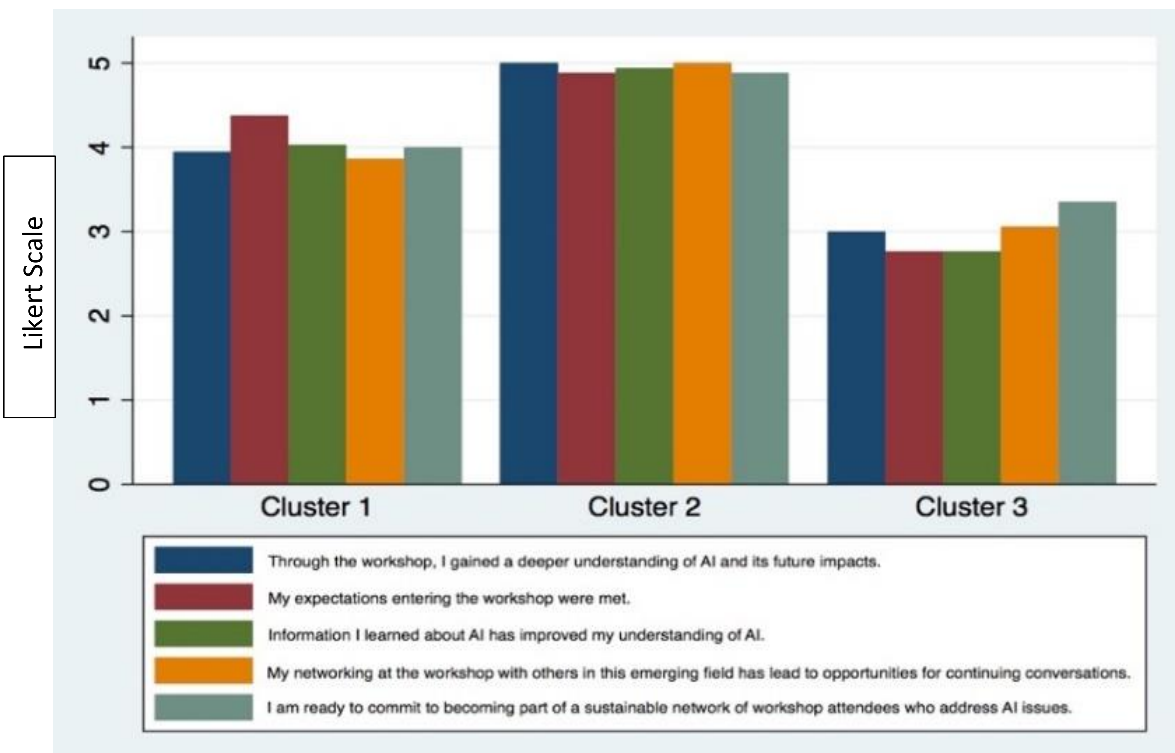

Figure 10: Cluster analysis on participants' satisfactions of AI workshop

Stem question: **Rate your level of satisfaction with each of the following aspects of the AI workshop:**

The cluster analysis revealed (see Figure 10) that the participants' expectations of the workshop can be categorized into three groups/clusters. We assigned descriptive, if unwieldy names, to these clusters by assigning high (mean > 4.5), middle (4.5 > mean > 3.5) and low (mean < 3.5) designations to Future Impact, met expectation, Understanding AI, Networking, Participating AI community. Table 3 shows the percentage of ethnicity/race (e.g., White, Black, and Asian) in each cluster.

Cluster 1 is labeled "Middle for all" to reflect satisfaction for Future Impact (mean=3.95 on a 5-point scale), Met expectation (mean=4.37), Understanding AI (mean=4.05), Networking (mean=3.84), and Participating AI community (mean =4.07). Cluster 2 is labeled "High for All" to reflect satisfaction for Future Impact (mean=5), Met expectation (mean=4.88), Understanding AI (mean=4.94), Networking (mean=5), and Participating AI community (mean =4.88). Cluster 3 is labeled "Low for all " and reflect satisfaction for Future Impact (mean=3.00), Met expectation (mean=2.76), Understanding AI (mean=2.76), Networking (mean=3.06), and Participating AI community (mean =3.35).

Table 3: Chi-Square Test of Independence (Cluster and Ethnicity) – Post Test

		Ethnicity/Race			
		White	Black	Asian	Total
Cluster	1	46.4%	33.3%	77.8%	55.8%
	2	17.9%	66.7%	5.6%	19.2%
	3	35.7%	0.0%	16.7%	25.0%
Total		100.0%	100.0%	100.0%	100.0%

Note: Sample size of cluster 1,2,3, and total are 29, 10, 13, and 52 respectively. Sample size of White, Black, Asian, and Total are 28, 6, 18, and 52 respectively.

Chi-squared tests of independence indicated that ethnicity/race reported in Table 3 is significantly different across the three clusters, $\chi 2$ (4) = 14.627, p = 0.006. White/Caucasian participants are the majority in this AI workshop and represent the total ratio of the clusters. Black/African-American participants are overrepresented in Cluster 2 (High for all) and underrepresented in Cluster 3 (low for all). Asian participants are overrepresented in Cluster 1 (Middle for all) and underrepresented Cluster 2 (High for all). According to chi-squared test of independence, Black/African American participants had more individuals who had higher satisfaction on all aspects (e.g., Future Impact, met expectation, Understanding AI, Networking, Participating AI community) of AI workshop. Meanwhile, more Asian participants fall into a profile of low satisfaction on all aspects. There was no significant association between other participants' characteristics (e.g., gender, job current setting, and year of experience) and the clusters.

Conclusion and Path Forward

The selection of the city of Pittsburgh as the location to host the Workshop on Artificial Intelligence and the Future of STEM and Societies became a fitting metaphor for how participants felt about what was needed to form lasting collaborations across industry, government, and academia. Rather than the gears metaphor (Figure 11) that was used at the start of the workshop, the building bridges metaphor (Figure 12) quickly became the prevailing way of thinking about the future of AI in STEM and Society, and the role of academia, industry, and government in shaping that future.

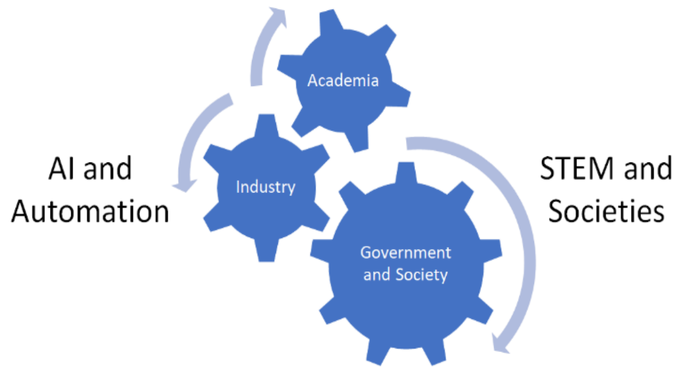

Figure 11: Two Themes of the Workshop: i) AI and Automation ii) STEM and Societies

Just take Pittsburgh, “the city of bridges.” Pittsburgh’s 446+ bridges connect the city (and within it, CMU), in ways that would not have been possible without them. The bridges metaphor also reminds us of the need to maintain connections, not just create them. There are countless examples of initiatives that were created by funding agencies attempting to bring industry, academia, and government together that originally started strong, but over time, were degraded, and in some cases, lost. The challenges facing society in the age of AI and automation are too significant for that to be the norm in the 21st century.

Figure 12: Pittsburgh and its Many Bridges

The 21st century is going to usher in significant changes in the way humans learn, work, and interact with technology. The diversity of people and topics represented at this workshop is a testament to the broad interest in AI, STEM and workforce development across different sectors and disciplines. As next steps, workshop participants outlined the following action items.

- Identify collaborative opportunities and set targets
- Maintain the network of stakeholders established as a result of this workshop
- Become change agents within our respective organizations in order to promote investments in STEM and AI

Workshop participants have already started to deliver on these action items including i) a submission of a National Science Foundation AI Institutes proposal, ii) discussion about ways to engage on the international stage through events such as the United Nations AI for Good Summit [20], and iii) discussion about the formation of a new Challenger Center in Pittsburgh. The Challenger Center is a STEM initiative that was established after the Challenger disaster, which seeks to engage students in space mission simulations and engineering and problem solving tasks [21]. In addition to the informal networking activities that occurred throughout the workshop, a formal contact list of workshop participants was created for participants who opted in and wanted their contact information shared with other participants.

Acknowledgements[1]

Funders/Sponsors

National Science Foundation
Mechanical Engineering, CMU

Any opinions, findings, or conclusions found in this report represent the collective views and opinions of the workshop participants and do not necessarily reflect the views of the sponsors or participating organizations.

Workshop Organizers and Advisers (Carnegie Mellon University)

Dr. Conrad Tucker, PI and Lead Workshop Organizer
Ms. Kelly McQuoid, Senior Director of Strategic Programs & Protocol, College of Engineering
Ms. Natasha Young, Administrative Coordinator, Mechanical Engineering
Dr. Allen Robinson, David and Susan Coulter Head and Raymond J. Lane Distinguished Professor of the Department of Mechanical Engineering
Dr. Jonathan Cagan, George Tallman and Florence Barrett Ladd Professor of Mechanical Engineering

Digital and Communications Team (Carnegie Mellon University)

Ms. Donna Sturgess, Chief Marketing Officer, College of Engineering
Mr. David Cochran, Senior Video Producer, College of Engineering
Ms. Bridget Decker, Senior Director, Marketing and Communications, College of Engineering
Mr. John Baldridge, Senior Manager, Digital Marketing and Communications, College of Engineering
Ms. Lisa Kulick, Senior Manager of Communications, Mechanical Engineering

Workshop Keynote Speakers and Panelists

Ms. Ursula Wynhoven, Representative to the United Nations in New York for the International Telecommunication Union (ITU)
Dr. James Garrett, Provost and Chief Academic Officer of Carnegie Mellon University
Ms. Mary Williams, Director of Outreach & Recruitment at the UNCF
General (Ret.) James Cartwright, former 8th Vice Chairman of the Joint Chiefs of Staff
Mr. Jeffrey Wilcox, Vice President for Digital Transformation at Lockheed Martin
Ms. Chau Dang, Product Manager at NVIDIA
Dr. Allen Robinson, David and Susan Coulter Head and Raymond J. Lane Distinguished Professor of the Department of Mechanical Engineering, CMU

[1] Listed in random order (not alphabetically) due to the reported biases in alphabetically-ordered listings https://researchonresearch.blog/2018/11/28/theres-lots-in-a-name/

Dr. Cathy Cavanaugh, Principal Program Manager for Professional Development in the education products group at Microsoft
Dr. Irene Petrick, Senior Director of Industrial Innovation in the Internet of Things Group at Intel
Dr. Guru Madhavan, Norman R. Augustine Senior Scholar and director of programs of the National Academy of Engineering
Dr. Arunan Skandarajah, Program Officer at the Bill and Melinda Gates Foundation on the Innovative Technology Solutions
Dr. Jonathan Cagan, George Tallman and Florence Barrett Ladd Professor of Mechanical Engineering, CMU
Colonel Robert Carter, Ph.D., Science & Technology Division, Directorate for Systems Integration, United States Army Futures Command
Dr. Barry Nagle, Senior Research Associate at the UNCF
Dr. Doug Barker, Chief Operating Officer at Energid Technologies
Dr. David Danks, Department Head & L.L. Thurstone Professor of Philosophy & Psychology, CMU
Mr. Chris Benson, Principal Artificial Intelligence Strategist at Lockheed Martin
Dr. Marsha Lovett, Associate Vice Provost for Teaching Innovation and Learning Analytics, Director of CMU's Eberly Center for Teaching Excellence & Educational Innovation and Teaching Professor in the Department of Psychology, CMU
Dr. Ben Amaba, Chief Technology Officer, Artificial Intelligence & Data Sciences for the Industrial & Engineering Sector at IBM
Ms. Bethanie Maples, Product Development Lead: education and machine learning at Google
Dr. Kiron Skinner, Director and Taube Professor, Institute for Politics and Strategy, CMU
Dr. Anthony Scriffignano, Senior Vice President, Chief Data Scientist at Dun & Bradstreet
Mr. Kelvin Harris, Director for Leadership Development and Alumni Relations at UNCF
Dr. Allyson Watson, Dean for the College of Education at the Florida Agricultural and Mechanical University
Dr. Marios Savvides, Bossa Nova Robotics Professor of Artificial Intelligence, Electrical & Computer Engineering, CMU
Dr. Andre Perry, Fellow in the Metropolitan Policy Program at Brookings
Dr. Akkarit Sangpetch, CMKM Program (Thailand) Director
Dr. Peter Engelke, Senior Fellow, Atlantic Council's Scowcroft Center on Strategy and Security
Dr. Patrick McSharry, Professor Department of Electrical and Computer Engineering, CMU
Mr. Nils Janus, Global Heading the AI activities of Covestro
Dr. David Bray, Senior Fellow, Institute for Human-Machine Cognition and PeaceTech Labs

Overview Talks on Federal Funding and Research Opportunities

Dr. Henry Kautz, NSF Division Director for Information & Intelligent Systems (IIS)
Dr. Kon-Well Wang, NSF Division Director of the Division of Engineering Education and Centers (EEC), Directorate of Engineering
Dr. Edward Berger, NSF Program Officer for Engineering Education Research in the Division of Engineering Education and Centers (Directorate for Engineering)
Dr. Grace C.Y. Peng, Director of Mathematical Modeling, Simulation and Analysis at the National Institute of Biomedical Imaging and Bioengineering (NIBIB) within the National Institutes of Health (NIH)
Dr. Thomas Fu, Director of the Advanced Naval Platforms Division of the Mission Capable, Persistent & Survivable Naval Platforms Department of the Office of Naval Research (ONR)
Dr. Tien Pham, Chief Scientist for the Computational & Information Sciences Directorate (CISD) at the U.S. Army Combat Capabilities Development Command (CCDC) Army Research Laboratory (ARL)

Workshop Panel Moderators

Dr. Rahul Telang, Professor and Ph.D. Program Chair at Heinz College, CMU
Dr. Eswaran Subrahmanian, Research Professor in the Engineering Research Accelerator and the Department of Engineering and Public Policy
Mr. Michael Krigsman, Industry analyst and host of CxOTalk (https://www.cxotalk.com)
Dr. Conrad Tucker, Arthur Hamerschlag Career Development Professor in Mechanical Engineering and Machine Learning (Courtesy) at CMU

NSF Survey Organizers and Workshop Data Analysts

Dr. Kathy Jackson, Affiliate Faculty, Department of Higher Education, Penn State
Dr. Jongho Park, Assistant Research Professor, Engineering Design, Penn State

Workshop Report Editors

Dr. Jonathan Cagan, George Tallman and Florence Barrett Ladd Professor of Mechanical Engineering, CMU
Dr. Kathy Jackson, Affiliate Faculty, Department of Higher Education, Penn State
Mr. Jacob Williamson-Rea, Content Writer, Engineering Research Accelerator, CMU
Dr. Darryl Farber, Assistant Teaching Professor, Engineering Design, Penn State

Workshop Breakout Session Leads and Scribes

1. Personalized Learning: Breakout Lead: **Dr. Amir Farimani**, Assistant Professor, Mechanical Engineering; CMU Student Volunteer(s): **Dule Shu**, **Yuxin Yao**, **Jamol Abdiev**
2. Broadening Access to STEM through AI: Breakout Lead: **Ms. Mary Williams**, Director of Outreach at the UNCF; CMU Student Volunteer(s): **Abishek Das**, **Japsimar Singh Wahi**
3. Data Ownership in the Age of AI: Breakout Lead: **Dr. Alessandro Acquisti**, Professor of Information Technology and Public Policy at the Heinz College; CMU Student Volunteer(s): **Malvika Singh**, **Aarushi Rai**
4. Ethics of AI in STEM and Society: Breakout Leads: **Mr. Chris Benson**, Principal Artificial Intelligence Strategist at Lockheed Martin; **Dr. Barry, Nagel**; CMU Student Volunteer(s): **Sakthi Prakash**, **Lihan Hu**, **Thomas Hurley**
4. The Future of Work and Lifelong Learning: Breakout Lead: **Dr. Darryl Farber**, Assistant Teaching Professor, **Engineering Design, Penn State; CMU Student Volunteer(s): Rounak Baheti, Venkata** Sanjay Renduchintala
6. The 21st Century "Classroom": Breakout Lead: **Dr. Chris McComb**, Assistant Professor of Engineering Design and Mechanical Engineering at Penn State; CMU Student Volunteer(s): **James Cunningham**, **Yiting Hui**

General Acknowledgements

In addition to the individuals and organizations acknowledged above, we would also like to extend a special thank you to each and every workshop participant for taking the time and effort to participate in the AI+STEM Workshop. We would also like to thank the staff who worked diligently behind the scenes to make the workshop a success.

Additional Workshop Resources and Links

Workshop Website: https://www.meche.engineering.cmu.edu/research/stem-ai-cmu.html
CMU Published News Article: https://engineering.cmu.edu/news-events/news/2019/12/09-ai-workshop.html
Appendix (including workshop survey questions and IRB information) can be found on the Official Workshop Website: https://www.meche.engineering.cmu.edu/research/stem-ai-cmu.html

References

[1] C. Greenhow, B. Robelia, and J. E. Hughes, “Learning, teaching, and scholarship in a digital age: Web 2.0 and classroom research: What path should we take now?,” *Educ. Res.*, vol. 38, no. 4, pp. 246–259, 2009.
[2] “Gates Millennium Scholars Program.” //www.gmsp.org/ (accessed Dec. 26, 2019).
[3] E. Rich, “Artificial intelligence and the humanities,” *Comput. Humanit.*, pp. 117–122, 1985.
[4] M. Storksdieck, “Critical information literacy as core skill for lifelong STEM learning in the 21st century: reflections on the desirability and feasibility for widespread science media education,” *Cult. Stud. Sci. Educ.*, vol. 11, no. 1, pp. 167–182, 2016.
[5] C. McNaught and P. Lam, “Using Wordle as a supplementary research tool.,” *Qual. Rep.*, vol. 15, no. 3, pp. 630–643, 2010.
[6] P. Y. Papalambros, “Gears and Transmissions: Quintessential Mechanical Design,” *J. Mech. Des.*, vol. 132, no. 6, 2010.
[7] “Low unemployment isn’t worth much if the jobs barely pay.” https://www.brookings.edu/blog/the-avenue/2020/01/08/low-unemployment-isnt-worth-much-if-the-jobs-barely-pay/ (accessed Mar. 03, 2020).
[8] A. Perry, J. Rothwell, and D. Harshbarger, “The Devaluation of Assets in Black Neighborhoods,” *Brook. Inst.*, 2018.
[9] A. Smith and J. Anderson, “AI, Robotics, and the Future of Jobs,” *Pew Res. Cent.*, vol. 6, 2014.
[10] R. Dillon, *The golden age of video games: The birth of a multibillion dollar industry*. AK Peters/CRC Press, 2016.
[11] S. Kiesler, R. E. Kraut, K. R. Koedinger, V. Aleven, and B. M. Mclaren, “Gamification in education: What, how, why bother?,” *Acad. Exch. Q.*, vol. 15, no. 2, pp. 1–5, 2011.
[12] “Definition of MUSCLE MEMORY.” https://www.merriam-webster.com/dictionary/muscle+memory (accessed Dec. 27, 2019).
[13] “I’d blush if I could: closing gender divides in digital skills through education - UNESCO Digital Library.” https://unesdoc.unesco.org/ark:/48223/pf0000367416.page=1 (accessed Mar. 03, 2020).
[14] J. Buolamwini and T. Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” in *Conference on Fairness, Accountability and Transparency*, Jan. 2018, pp. 77–91, Accessed: Dec. 27, 2019. [Online]. Available: http://proceedings.mlr.press/v81/buolamwini18a.html.
[15] I. Goodfellow *et al.*, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
[16] M. Miller, J. P. De Clerck, S. A. Sorby, L. M. Roberts, W. J. Endres, and K. D. Hale, “Meeting the NAE Grand Challenge: Personalized Learning for Engineering Students through Instruction on Megacognition and Motivation Strategies,” in *120th ASEE Annual Conference Proceedings*, 2013.
[17] “Definition of DATA.” https://www.merriam-webster.com/dictionary/data (accessed Dec. 24, 2019).
[18] H.-Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. Freeman, “Eulerian video magnification for revealing subtle changes in the world,” *ACM Trans. Graph. TOG*, vol. 31, no. 4, pp. 1–8, 2012.
[19] E. Politou, E. Alepis, and C. Patsakis, “Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions,” *J. Cybersecurity*, vol. 4, no. 1, p. tyy001, 2018.
[20] “Home - AI for Good Global Summit.” https://aiforgood.itu.int/ (accessed Mar. 31, 2020).
[21] “Home - Challenger Center.” https://www.challenger.org/ (accessed Mar. 29, 2020).