Carnegie Mellon University

Statement to the US Senate

By
Jodi Forlizzi
Herbert A. Simon Professor
Human Computer Interaction Institute, School of Computer Science
Faculty Lead, Responsible AI, Block Center for Technology and Society
Carnegie Mellon University

AI Insight Forum on AI Innovation
Tuesday, October 24, 2023

Majority Leader Schumer, Senators Rounds, Heinrich, and Young, thank you for hosting this important event today, and for giving me the opportunity to share this statement.

My name is Jodi Forlizzi and I am the Herbert A. Simon Professor in Human-Computer Interaction in the School of Computer Science at Carnegie Mellon University. I also serve as a Responsible AI Faculty Lead in the Block Center for Technology and Society and the Associate Dean of Diversity, Equity, and Inclusion in the School of Computer Science. My perspective is shaped by my background of over 20 years of work in design, in studying and developing responsible human-AI interaction in the product development process, and in understanding how AI affects work, primarily the work of front-line, face-to-face service workers. My work has demonstrated that strategies and methodologies that enhance worker engagement in the development of new technologies can contribute to ensuring that the US remains a leader in the adoption of AI innovations.

Background on AI Innovation

Discussions about AI hold a lot of promise, but also a lot of hype. A 2022 survey by TechRepublic [1] reported that 97% of CEOs planned a significant investment in AI.

Generative AI now offers the ability to move fluidly between text and image prompts [2]; AI is being used to advance healthcare through prediction and prevention of disease [3]; to reduce recidivism rates for those needing mental health services [4]; and to improve educational outcomes for students who are at risk of not graduating [5].

The massive success of AI innovations implies a robust, responsible AI innovation process. However, a closer examination of AI innovation reveals this to not be true. Nearly 85% of AI innovations within companies fail before deployment, and of those that are deployed, nearly 40% fail to be successful [6].

There are four main reasons this happens:

- **Model Performance:** Teams cannot achieve the model performance they need to be "good enough." This might be because the problem is too hard or this can be for internal reasons, such as the group with access to the data will never provide the needed access. A good example of this is the Galactica LLM released by Meta. It was supposed to answer scientific questions. It had to be turned off after two days.
- **Service Value:** The system will never generate enough value to make it worthwhile for the service provider. Teams often never consider what it will cost to build or to operate their project. Or they fail to test an underlying assumption about how something will generate revenue. Amazon Alexa provides a nice example of this. Amazon subsidized the hardware cost and achieved a great uptake of the Alexa devices. However, the revenue model was based on the assumption people would shop with Alexa. It does not seem like this assumption was tested.
- **Customer Value:** The system does not generate value for the user, so the user never adopts or uses the system. For example, most AI systems developed for healthcare fall into this failure space. They don't produce value for clinicians, so clinicians don't use them.
- **Ethical Issues:** The system has issues with data privacy, algorithmic bias, or some other type of unintended harm that creates risk for the service provider. These problems are often revealed during deployment, when it is too late to correct course. The call to develop responsible AI continues to grow.

AI Innovation and Workers

Along with UNITE HERE!, the largest hospitality union in the US, which is part of the AFL-CIO, I have secured funding from the National Science Foundation to produce insightful research on how AI should be designed, developed, and deployed in the hospitality industry. This work reflects the charge included in the CHIPs Act for NSF's

Technology, Innovation, and Partnerships directorate to help support mutually beneficial partnerships with labor organizations that can support use-inspired translational research (Subtitle G Section 10383 (3)). This project emerged from a broader engagement with the AFL-CIO's Technology Institute which is committed to supporting technology innovation by engaging workers in the innovation process. The project highlights that this model can contribute to addressing critical challenges impacting the adoption of AI and open new pathways for industry collaboration.

When a guest stays in a union hotel, the service staff are members of UNITE HERE. Workers at many of these properties, which can be owned, managed, and branded by up to three different organizations, are increasingly encountering AI in their day to day work. For example, algorithmic managers assist housekeepers and managers to manage operations and assist in the workflow of cleaning rooms. They also allow for customized room assignments, to deal with issues like getting ready for conference attendees and prioritizing the cleaning of rooms where guests have checked out.

Housekeeping is a physically demanding job, largely performed by middle-aged women of color and immigrants. Many do not speak English, or possess much digital literacy. However, they take pride in the work that they do and the face-to-face interactions they have with their customers. Many have been doing their jobs for two or three decades.

However, AI can be problematic in this space, because the needs of workers, customers, and managers are prioritized differently. And workers are at the bottom of this list. For housekeepers, algorithmic managers (AMs) increase work, increase job requirements, and decrease worker autonomy. Instead of letting housekeepers clean rooms in the order that makes the most sense to them based on their ability to complete their room quotas with a minimum of wear and tear on their bodies, AMs send them back and forth and up and down in hallways and elevators, while they push 200-300 pound carts. We have heard again and again from housekeepers that the AM "wastes my time." AMs increase wear on the worker by assigning several check-out rooms, which require heavy cleaning, back to back as opposed to alternating them with the lighter physical requirements of rooms in which only sheets and towels need changing. Housekeeping is also an entry level job that traditionally did not require technology skills or even fluent English. This, combined with typical connectivity issues, has altered the job of the housekeeper greatly, with little to no increased training or increased compensation.

Sustainable AI Innovation for Workers

While the entire lifecycle of AI innovation – scoping, procurement, designing, testing, deploying, and using – needs to have guidelines and practices in place to maximize benefit and minimize risks. Some efforts are underway, such as those at Apple [7], Google [8], and Microsoft [9]. Researchers have also created frameworks and other mechanisms for exploring ethics and developing responsible AI systems. One example of this is the NIST Risk Management Framework [10]. In a policy brief published by the Responsible AI Initiative at the Block Center at Carnegie Mellon University, we highlighted some of the unique challenges of AI Accountability and laid out a set of policy recommendations [11].

Today I want to focus on a few recommendations to help think about how to maximize investment in sustainable AI innovation.

1. **Bring workers into the design, development, training, deployment process**

My experience working with UNITE HERE! reinforced the importance of building teams to understand work and how it is affected by new technology. As we explored how hotel guest room attendants use algorithmic managers to move from room to room during their workday, conversations with the attendants themselves revealed one perspective about the work. Talking to their managers who did the room assignments revealed another perspective. Talking to hotel management and the manufacturer of the software revealed additional, sometimes conflicting, perspectives. We need to develop processes to collaborate with workers, planners, and decision makers to surface biases and mitigate them as we seek to achieve better design, development, and implementation around AI and automation in the hospitality industry.

Too many times, the response is to choose the simplest technology solution to create rapid improvement. Such a narrow view can result in poor solutions that cost a lot of money and perpetuate the problem. Managerial strategies, particularly whether a firm pursues “high-road” or “low-road” management practices regarding labor, are crucial determinants for worker outcomes [12; 13]. High road practices compete on quality of service, achieved through investment in human capital and an innovative mindset, generating jobs where workers have agency and voice [14; 12]. Low road practices implement technology on the basis of cost and price minimization.

For example, early automated checkouts in grocery stores were created without rich insights on how staff and customers might interact with these automated systems. Traditional staff positions were reduced, while the new systems created the need for a new kind of role overseeing interaction with the new terminals. The net result was reduced efficiency in the checkout lines and the need for additional skilled workers. In

contrast to this low-road automation example, in high-road implementation, worker well-being on the job would be a concern, with a focus on on-the-job success, workplaces that support worker success, job flexibility and adequate pay, and a resulting pool of talented and productive workers that would perform high-quality, revenue-generating work. Thoughtful workforce policy can assert the importance of a systemic view of work in the face of new technology implementation.

2. **Develop a process for developing sustainable AI, figure out what is generalizable and what is idiosyncratic**

There is a need for federal research to develop processes for designing, developing, and deploying sustainable AI for workers. This could likely be sector focused, and idiosyncrasies will always exist. For example, my previous research studying a service robot in a hospital showed how in the same building, in two wings with identical structure, where the robot was programmed to do the exact same job, the service the robot provided was perceived by workers very differently in the two different wings. This is because the work in each wing – cancer patients and postpartum patients – was different in nature and created different impressions of, and interactions with, the technology [15].

Processes for developing sustainable AI would ideally help to articulate different perspectives and mitigate the biases that naturally arise in working with these kinds of teams. Professionals and managers may feel that they have more expertise than any AI system could offer, and may never want to use AI. Workers may feel threatened by AI, and may also avoid it. The goal is to both produce concepts that span the entire problem/opportunity space and that sit at the intersection of AI's capabilities and stakeholder needs. A second goal is to identify opportunity and risks at the earliest stages of an innovation process in order to choose solutions that are more likely to be successful. A third goal is to develop metrics and measures to rate the overall sustainability of potential solutions.

3. **Design sustainable AI from the very beginning of the product development process**

Finally, we need processes and methods to help product managers and development teams to become aware of potential risks early on in the design and development process, not after the product has been built. Research suggests that teams are interested in avoiding fairness, accountability, transparency and ethics, or FATE problems, in the first place [16], and that the most effective way to build more sustainable AI is to address responsibility during ideation and problem formulation [17]. This is a change from current processes that unintentionally focus on high-risk projects

and that only consider responsible AI after problem formulation. To assist in this process, frameworks must be created that help product teams articulate the value proposition of their envisioned product, for whom that value is generated (i.e., data beneficiaries), whose data is processed to unlock that value (i.e., data subjects), who can be impacted by the data pipeline (i.e., data victims), and by which AI risk (e.g., surveillance).

4. **Make digital literacy training programs readily available**

In our initial study of the use of housekeeping AMs by employees, supervisors, and managers, we learned that digital literacy is a key factor in the success of this technology. We visited typical hospitality training facilities, talked to trainers, and reviewed digital literacy curricula. Currently, training on digital literacy is fairly sparse for many hotel employees. While existing digital literacy programs are successful in getting people to their first jobs, they often fall short of serving workers over long careers, and they can be insufficient for workers who are not digital natives.

Funding is needed to increase access to, and competency with, digital technologies for workers of all kinds, but particularly older workers, by developing, deploying, and evaluating supplemental digital literacy training programs to serve these individuals. Our research has shown that some workers struggle with digital literacy, retaining only very basic skills. Our research has revealed opportunities to provide additional training to address the needs of these workers and to improve their day-to-day work. Assistance is needed for these workers to more fully understand and utilize technology and remain in the workforce. We can evolve with a set of key educational concepts around digital literacy and AI that can serve a diverse workforce. Now is the time to upskill every worker in every workplace.

Thank you again for convening this vital session. As an academic who has spent their career researching ways to ensure that breakthroughs in technology empower individuals and communities I appreciate the leadership that you are bringing to developing policies that can ensure that the US is the global leader in responsible and trustworthy AI innovation.

---

[1] Bernard, Allen. “U.S. businesses investing in artificial intelligence this year to improve efficiency and save money.” TechRepublic: Artificial Intelligence. June 4, 2021. Accessed

on October 14, 2022:
https://www.techrepublic.com/article/ai-gaining-traction-in-the-workplace/

[2] Madhur Garg. (2023). Meat AI Introduces AnyMAL: The Future of Multimodal Language Models Bridging Text, Images, Videos, Audio, and Motion Sensor Data. October 5, 2023. Accessed on October 14, 2022:
https://www.marktechpost.com/2023/10/05/meta-ai-introduces-anymal-the-future-of-multimodal-language-models-bridging-text-images-videos-audio-and-motion-sensor-data/

[3] Kilic, Arman, Daniel Dochtermann, Rema Padman, James K. Miller, and Artur Dubrawski. "Using machine learning to improve risk prediction in durable left ventricular assist devices." PLoS One 16, no. 3 (2021): e0247866.

[4] Bauman, Matthew J., Kate S. Boxer, Tzu-Yun Lin, Erika Salomon, Hareem Naveed, Lauren Haynes, Joe Walsh et al. "Reducing incarceration through prioritized interventions." In Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, pp. 1-8. 2018.

[5] Lakkaraju, Himabindu, Everaldo Aguiar, Carl Shan, David Miller, Nasir Bhanpuri, Rayid Ghani, and Kecia L. Addison. "A machine learning framework to identify students at risk of adverse academic outcomes." In Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining, pp. 1909-1918. 2015.

[6] Weiner, Joyce. “Amazon Is Said to Plan to Lay Off Thousands of Employees.” The New York Times. November 14, 2022. Accessed on Nov 22, 2022.
https://www.nytimes.com/2022/11/14/technology/amazon-layoffs.html

[7] Apple. 2019. Human Interface Guidelines: Machine Learning. Accessed on Nov. 18, 2022:
https://developer.apple.com/design/human-interface-guidelines/technologies/machine-learning/introduction/

[8] Google PAIR. 2019. People + AI Guidebook. Accessed on Nov. 22, 2022:
http://pair.withgoogle.com/guidebook23

[9] Amershi, Saleema, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh et al. "Guidelines for human-AI interaction." In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, pp. 1-13. 2019.

[10] NIST Risk Management Framework (2022). https://csrc.nist.gov/projects/risk-management/about-rmf, accessed July 1, 2023.

[11] Toward AI Accountability: Policy Ideas for Moving Beyond a Self-Regulatory Approach. January, 2023. https://www.cmu.edu/block-center/responsible-ai/cmu_blockcenter_rai-memo_final.pdf, accessed October 1, 2023.

[12] Boushey, Heather, & Rinz, Kevin. (2022, April 6). Blocking the low road and paving the high road: Management practices to improve productivity. [Issue Brief] Washington, DC: The White House. https://www.whitehouse.gov/cea/written-materials/2022/04/06/blocking-the-low-road-and-paving-the-high-road-management-practices-to-improve-productivity/.

[13] Institute for the Future of Work (IFOW). (2022, May). Case for importance: Understanding the impacts of technology adoption on 'good work.' IFOW. https://www.ifow.org/publications/impacts-technology-adoption-work

[14] Kresge, Lisa. (2020a, November). Data and algorithms in the workplace: A primer on new technologies. Berkeley, CA: University of California Labor Center. https://laborcenter.berkeley.edu/working-paper-data-and-algorithms-in-the-workplace-a-primer-on-new-technologies/

[15] Mutlu, Bilge, and Jodi Forlizzi. "Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction." In Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction, pp. 287-294. 2008.

[16] Deng, Wesley Hanwen, Bill Boyuan Guo, Alicia Devos, Hong Shen, Motahhare Eslami, and Kenneth Holstein. "Understanding Practices, Challenges, and Opportunities for User-Driven Algorithm Auditing in Industry Practice." arXiv preprint arXiv:2210.03709 (2022).

[17] Holstein, Kenneth, Jennifer Wortman Vaughan, Hal Daumé III, Miro Dudik, and Hanna Wallach. "Improving fairness in machine learning systems: What do industry practitioners need?." In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, pp. 1-16. 2019.