

PACE: Participatory AI for Community Engagement

Saad Hassan¹, Syeda Mah Noor Asad¹, Motahhare Eslami², Nicholas Mattei¹,
Aron Culotta¹, John Zimmerman²

¹Tulane University, New Orleans, LA

²Carnegie Mellon University, Pittsburgh, PA

saadhassan@tulane.edu, sasad1@tulane.edu, meslami@andrew.cmu.edu, nsmattei@tulane.edu,
aculotta@tulane.edu, johnz@andrew.cmu.edu

Abstract

Public sector leverages artificial intelligence (AI) to enhance the efficiency, transparency, and accountability of civic operations and public services. This includes initiatives such as predictive waste management, facial recognition for identification, and advanced tools in the criminal justice system. While public-sector AI can improve efficiency and accountability, it also has the potential to perpetuate biases, infringe on privacy, and marginalize vulnerable groups. Responsible AI (RAI) research aims to address these concerns by focusing on fairness and equity through participatory AI. We invite researchers, community members, and public sector workers to collaborate on designing, developing, and deploying RAI systems that enhance public sector accountability and transparency. Key topics include raising awareness of AI's impact on the public sector, improving access to AI auditing tools, building public engagement capacity, fostering early community involvement to align AI innovations with public needs, and promoting accessible and inclusive participation in AI development. The workshop will feature a keynote, a short paper session, and two discussion-oriented activities. Our goal is to create a platform for exchanging ideas and developing strategies to design community-engaged RAI systems while mitigating the potential harms of AI and maximizing its benefits in the public sector.

Background and Introduction

Today, Artificial Intelligence (AI) transforms how citizens and cities interact. It impacts public sector decision making, e.g., criminal justice risk scores (Goel, Rao, and Shroff 2016), resource allocation (Mehr, Ash, and Fellow 2017), and offers community groups ways to monitor and audit the actions of their governments, e.g., community-led AI tools to identify biases in criminal justice systems. AI's public sector impact is profound, yet dual-sided. On the positive side, it makes cities smarter, such as by predicting waste management demands to improve quality of life (Sroka 2023), or by streamlining bureaucratic processes to make operations more efficient. It also enables communities to hold the government more accountable. For example, communities in Chicago use participatory AI to analyze police data to identify potential biases and misconduct (BTSurface 2024). On the negative side, AI can reinforce and even

magnify human biases, e.g., racial bias in predictive policing systems (Alikhademi et al. 2022), invade citizen privacy, e.g., through enhanced surveillance (Eubanks 2018), automate work that leads to job displacement, e.g., AI replacing public sector workers in transportation for roadway inspections (Pittman 2022), and silence marginalized voices (e.g., online systems that misclassify disability-related content as negative sentiment or even hate speech (Hassan, Huenerfauth, and Alm 2021; Hutchinson et al. 2020; Morris 2020)).

Concerns about AI harms gave rise to Responsible AI (RAI) research, work committed to creating new AI innovation approaches meant to increase fairness and equity—traditionally peripheral topics in computing research (Eubanks 2018). Recently, the RAI research community has begun to engage with the public sector, driven by the public sector's commitment to equity, fairness, and societal improvement (Boyne 2003). Unlike the commercial sector, which primarily measures success through profit-driven metrics like sales, usage, and cost reduction, the public sector strives to balance efficiency with a strong dedication to equity, fairness, and societal advancement. Emerging RAI research emphasizes collaborative efforts between communities and municipalities to assess and improve public sector accountability and transparency through several paths including a) investigating the use of participatory design (PD) to include impacted, marginalized communities in designing public sector AI systems (Zytka et al. 2022) and b) utilizing auditing to assess quality and find biases within public service decision making (Loi and Spielkamp 2021).

RAI research has made preliminary advances in using PD to engage communities in public sector AI innovation and in developing auditing tools that reveal ethical problems. However, this research faces several challenges, limiting its impact. (a) *Awareness Challenge*: Community members may not be aware of, or comprehend the impact of public sector innovations. For example, recent work by organizer Eslami et al. revealed a strikingly low level of public awareness (below 10%) regarding a city's use of AI. This limits the participation of impacted, marginalized communities in the process of design, development, and deployment of public sector AI (Lin et al. 2023). (b) *Access Challenge*: People subject to public-sector decision-making processes often lack access or understanding of the process. They have no way to interrogate or audit these decisions. For instance,

an AI child welfare risk assessment tool adopted by several US cities influenced case worker decisions. These negatively impacted those who had no way to examine AI's influence (Stapleton et al. 2022). (c) *Ability Challenge*: Auditing tools are typically built and used by a small number of experts, often system developers, data scientists, or a select group of users who have the technical background to provide feedback. This prevents members of the public from meaningfully engaging in the creation of AI systems that monitor and improve city functioning. (d) *Alignment Challenge*: Tools and processes built to audit public-sector decision-making typically get created at the very end of the design/develop/deploy lifecycle. Innovation teams overlook opportunities to engage community members in the early stages of design to align innovation goals with the needs of the community. (e) *Accessibility Challenge*: Finally, the lack of inclusive channels for engagement prevents diverse groups with varied abilities and backgrounds from fully contributing to the development of public sector AI systems as well as utilizing AI systems in monitoring public sector practices.

Our workshop aims to discuss these challenges and opportunities to mitigate such challenges. Further, we will ideate advancing RAI by advancing PD and increasing the role of auditing in the public sector. We invite researchers working with communities, particularly disadvantaged and marginalized populations, and public sector workers, to engage in discussions on the design, development, and deployment of public sector AI systems. The goal is to foster discussions around AI innovation approaches for the public sector that maximize benefits and reduce harms.

Organizers

1. Motahhare Eslami's (Assistant Professor, Human-Computer Interaction Institute) research goal is to investigate the existing accountability challenges such as auditing in algorithmic systems and to empower the users of algorithmic systems, particularly those who belong to marginalized communities and those whose decisions impact marginalized communities, make transparent, fair, and informed decisions in interaction with algorithmic systems. Motahhare is named one of the "100 Brilliant Women in AI Ethics" and her work has been covered in mainstream media such as Time, The Washington Post, Huffington Post, the BBC, Fortune, Quartz, and WIRED.
2. John Zimmerman (Professor, Artificial Intelligence and Human-Computer Interaction) is an interaction designer and researcher who teaches courses on interaction design, service innovation, design theory, and HCI methods. His research covers four areas: blending human and machine intelligence, designing emotionally engaging digital products, innovating public services through social computing, and integrating design inquiry with scientific and engineering methods. His notable projects include a clinician decision support system, an alarm clock for parents, the Tiramisu transit app, and the book "Design Research through Practice." Before the HCI Institute, John researched personalized TV systems at Philips

Research.

3. Saad Hassan (Assistant Professor, Computer Science) researches accessible computing and HCI, with a focus on AI-based technologies to promote inclusion, facilitate learning, and encourage creative expression among individuals with disabilities. Saad has worked with several industry partners, including Google Research and Meta Reality Labs and community partners, including Deaf Professional Arts Network (DPAN) and IPSOS. He teaches courses on HCI and Data Science.
4. Aron Culotta (Associate Professor, Computer Science) conducts research on natural language processing, machine learning, and artificial intelligence. As director of Tulane's Center for Community Engaged AI, he has expertise in building AI systems in partnership with community groups, and in designing and evaluating NLP systems. He teaches courses on AI, NLP, and Data Science.
5. Nicholas Mattei (Associate Professor, Computer Science) conducts research on artificial intelligence and its applications to theoretical and applied decision making. He is the co-director of Tulane's Center for Community Engaged AI and works with community groups and the City of New Orleans on open data projects. He teaches courses on AI, Data Science, and the impact of these technologies on society.
6. Syeda Mah Noor Asad (PhD Student, Computer Science) is a student at Tulane University. Her research focuses on health technologies, computational linguistics, and HCI.

Workshop Announcement

AI is transforming interactions between citizens and cities, revolutionizing civic infrastructure and public-sector decision-making. While AI enhances urban efficiency and accountability, it also risks perpetuating biases, infringing on privacy, and marginalizing vulnerable groups, necessitating responsible design and continuous oversight. Although researchers recommend participatory design, there is little guidance on gathering and incorporating feedback from community stakeholders in the public sector.

Our workshop will bring together researchers and practitioners focused on fairness and equity in public-sector AI through participatory design. Topics of interest include approaches and tools for raising awareness about public-sector AI, fostering alignment through early-stage community involvement, building capacity for public engagement in AI design and auditing, enhancing access to civic decisions through social and collaborative tools, and promoting accessibility for diverse participation in AI development.

We invite submissions of 5000 words or less (excluding references and appendices) for work-in-progress papers, position papers, or artifact papers. Examples include literature reviews on shortcomings or success stories in public-sector AI, proposed new solutions for the public sector, or community-engaged auditing tools for existing systems, etc. Our workshop will also feature a keynote and two discussion activities on participatory AI and accessibility. For more details visit: <https://tulaneecs.github.io/PACE/>

Time Slot	Theme
09:00 AM - 09:15 AM	Opening Remarks
09:15 AM - 10:00 AM	Keynote
10:15 AM - 11:15 PM	Paper Session
11:15 PM - 12:00 PM	Activity 1
12:00 PM - 01:00 PM	Lunch Break
01:00 PM - 01:45 PM	Activity 2
01:45 PM - 02:00 PM	Closing Remarks

Table 1: Tentative plan for the workshop.

Workshop Plan

Our workshop will feature opening and closing remarks from the organizers, a keynote, a paper session, and two discussion-oriented activities. Table 1 outlines our tentative plan. The final plan will be posted on the workshop website by the end of August.

Keynotes

Our keynote speaker will be from a community organization working on public sector AI based on our collaboration with several community organizations. This includes the Pittsburgh Task Force on Public Algorithms, Community Forge, Court Watch NOLA, Vera Institute of Justice, Eye on Surveillance, The Data Center New Orleans, Deaf Eye Consulting, and Center for Planning Excellence Louisiana (CPEX).

Paper Sessions

We will accept between 6-8 papers, allowing approximately 8-10 minutes per paper presentation and following Q&A. We will encourage the authors to make the talks as engaging as possible. The paper session will be chaired by one of the organizers.

Authors will be asked to share their presentation slides a week before the workshop date.

We will invite attendees, including authors of unsuccessful paper submissions, to bring posters to the venue.

Discussion Activities

We plan to organize two discussion activities during the workshops. Depending on the number of attendees, we will divide participants into multiple groups based on their interests in each activity. At the end of each group discussion, we will ask a group member to summarize the discussion for all the attendees. Examples of possible discussion activities include:

- **Activity 1: Designing Auditable Public-sector AI**

- **Discussion Focus:** How can the public sector innovate with AI by aligning community needs and AI capabilities, centering the concept of public auditing of municipal AI systems across all stages?
- **Potential Topics:** Understanding how to best engage impacted stakeholders, including city managers, public sector workers, and citizens, in a participatory manner. Emphasizing the concept of collective auditing

from the beginning, addressing how the public can audit systems that handle sensitive and private citizen data. Designing collaborative auditing tools for public sector AI.

- **Activity 2: Accessibility in Public Sector AI**

- **Discussion Focus:** How can we ensure that participatory design approaches to RAI include people with disabilities, who are often overlooked despite existing across all community groups?
- **Potential Topics:** Understanding how to guide the inclusion of people with visible and invisible disabilities in PD, prototyping, and collaborative auditing activities identified in the previous discussion points. Assess stakeholder willingness to address a wide range of needs throughout the AI innovation process.

Logistics

The workshop will accommodate up to 50 participants and will require standard audio-visual equipment, including two to three microphones for presentations and keynotes, as well as high-speed internet for remote speakers. Accessibility requirements will be communicated to the chairs after the paper selection process and finalization of the event plan.

Advertising

The workshop will be advertised through email distribution lists of relevant conferences and research labs in academia and industry. We will also use social conference platforms such as Whova to advertise the workshop to several venues, including but not limited to ACM SIGCHI, ACL, and AAAI venues. We will also set up a social media page and post summaries of keynotes and discussions.

We plan on setting up bi-weekly meetings to discuss the tasks with the team. We will formulate a program committee (PC) with experts in the fields of RAI and HCI by August. They will assist us in curating the workshop by sharing the calls for papers in their institutions and helping with the reviewing process for the submissions.

Participants who are interested in being part of the workshop will submit 3-5 pages short papers (excluding references). We will set up a reviewing management system for the submissions. Currently, we have identified two such systems: EasyChair and OpenReview. Two of the organizers with extensive prior experience in using these systems will lead the reviews and discussion process.

Website

We plan on setting up a website to advertise and disseminate the call for participation and other workshop information. The website will also be used to share the findings of the workshop, and to support future engagements with the attendees and the community. The website will be maintained by organizer Nicholas Mattei.

Submissions and Reviewing

Our workshop accepts submissions in the following categories:

- **Work-in-progress papers:** These submissions can present initial research findings or innovative ideas that are in the early stages of development. They should highlight the next steps for further research.
- **Position papers:** These can consist of literature reviews that reveal systematic shortcomings in the way researchers have been thinking about community-engaged AI. They can also propose new RAI solutions for the public sector, with or without (light) evaluations.
- **Artifact papers:** These could be in the form of software, prototypes, pictorials, etc. The artifacts must be described in an accompanying short 3-5 page paper.

The submissions should not be longer than 5,000 words (on average 3-5 pages using a double-column AAAI article template). Submissions should be in accessible PDF format. References and appendices do not count toward word limits.

We will use an anonymous review process in which the authors must anonymize their submissions as well as any artifacts or supplementary material separately shared. The submissions will be peer-reviewed by at least 2 reviewers who are experts in the field. Reviews for all the submissions will be shared with the authors. Notification of acceptance will be released by the first week of September. Accepted submissions will be uploaded on the website. At least one author of each accepted position paper or workshop artifact must attend the workshop and all participants must register for the workshop.

Workshop Accessibility Considerations

We will ensure the workshop is accessible to all attendees, including those who are blind, have low vision, or are d/Deaf and hard of hearing. We will work with conference organizers to ensure interpreter support. Submissions and presentations will follow accessibility guidelines, including alt text for images.

Accessibility will be integrated into all planning and reviewing stages. Participants can indicate specific requirements beforehand, and an accessibility contact will be provided on our website.

Post Workshop Plans

After the workshop, we plan to write a concise position paper summarizing the research findings from the accepted submissions, keynote, and ideas from the two discussion activities. All accepted papers will be available on the website for at least a year. Additionally, we will create a Discord server for all authors and attendees to facilitate ongoing discussions post-workshop.

Formatted bibliographies should look like the following examples. You should use BibTeX to generate the references. Missing fields are unacceptable when compiling references, and usually indicate that you are using the wrong type of entry (BibTeX class).

References

- Alikhademi, K.; Drobina, E.; Prioleau, D.; Richardson, B.; Purves, D.; and Gilbert, J. E. 2022. Correction to: A review of predictive policing from the perspective of fairness. *Artif. Intell. Law*, 30(1): 19–20.
- Boyne, G. A. 2003. Sources of Public Service Improvement: A Critical Review and Research Agenda. *J Public Adm Res Theory*, 13(3): 367–394.
- BTSurface. 2024. Investigations. <https://www.btsurface.com/investigations>. Accessed: 2024-07-25.
- Eubanks, V. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Publishing Group.
- Goel, S.; Rao, J. M.; and Shroff, R. 2016. Personalized Risk Assessments in the Criminal Justice System. *Am. Econ. Rev.*, 106(5): 119–123.
- Hassan, S.; Huenerfauth, M.; and Alm, C. O. 2021. Unpacking the interdependent systems of discrimination: Ableist bias in NLP systems through an intersectional lens. *arXiv preprint arXiv:2110.00521*.
- Hutchinson, B.; Prabhakaran, V.; Denton, E.; Webster, K.; Zhong, Y.; and Denuyl, S. 2020. Social Biases in NLP Models as Barriers for Persons with Disabilities.
- Lin, Y.-R.; Schwanke, B.; Farzan, R.; Fan, B.; Eslami, M.; Shen, H.; and Fox, S. 2023. The Public Algorithms Survey in Allegheny County.
- Loi, M.; and Spielkamp, M. 2021. Towards Accountability in the Use of Artificial Intelligence for Public Administrations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, 757–766. New York, NY, USA: Association for Computing Machinery.
- Mehr, H.; Ash, H.; and Fellow, D. 2017. Artificial intelligence for citizen services and government. *Ash Cent. Democr. Gov. Innov. Harvard Kennedy Sch.*, no. August, 1–12.
- Morris, M. R. 2020. AI and accessibility. *Commun. ACM*, 63(6): 35–37.
- Pittman, M. 2022. AI And The Future Of Government Work. *Forbes Magazine*.
- Sroka, N. 2023. How AI is revolutionizing Solid Waste Management. Accessed: 2024-4-19.
- Stapleton, L.; Lee, M. H.; Qing, D.; Wright, M.; Chouldechova, A.; Holstein, K.; Wu, Z. S.; and Zhu, H. 2022. Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 1162–1177. New York, NY, USA: Association for Computing Machinery.
- Zytka, D.; J. Wisniewski, P.; Guha, S.; P. S. Baumer, E.; and Lee, M. K. 2022. Participatory Design of AI Systems: Opportunities and Challenges Across Diverse Users, Relationships, and Application Domains. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, number Article 154 in CHI EA '22, 1–4. New York, NY, USA: Association for Computing Machinery.