

Human Interactions with Artificial Intelligences: Research Agendas for Education, Health, Creative Expression, and Governance

Brian Southwell, Lee Rainie, Katie Bowler Young, Maria Gallardo-Williams, Tony Kipkemboi, Jim Redden, Matthew Strobl, Jon Accarrino, Evan Bunnage, Robert Chew, Schuyler DeBree, Heather Griffiths, Haley Hickman, Rachel Levy, Selena Monk, Rachel Page, Jessica Reif, Rebecca Ivic, and Elizabeth Mannshardt



RTI Press publication OP-0100-2603

RTI International is an independent scientific research institute dedicated to improving the human condition. The RTI Press mission is to disseminate information about RTI research, analytic tools, and technical expertise to a global audience. RTI Press publications are peer-reviewed by at least two independent substantive experts and one or more Press editors.

Suggested Citation

Southwell, B., Rainie, L., Bowler Young, K., Gallardo-Williams, M., Kipkemboi, T., Redden, J., Strobl, M., Accarrino, J., Bunnage, E., Chew, R., DeBree, S., Griffiths, H., Hickman, H., Levy, R., Monk, S., Page, R., Reif, J., Ivic, R., and Mannshardt, E. (2026). *Human interactions with artificial intelligences: Research agendas for education, health, creative expression, and governance*. RTI Press Publication No. OP-0100-2603. RTI Press. <https://doi.org/10.3768/rtipress.2026.op.0100.2603>

Cover photo: Siobahn Day Grady, Melissa Hodge-Penn, and Rachel (Ray) Levy at the September 2025 summit entitled The Human Edge: Our Future with Artificial Intelligences in Research Triangle Park, NC. Photo credit: Markeith Gentry, RTI International.

This publication is part of the RTI Press Occasional Paper series.

RTI International
3040 East Cornwallis Road
Durham, NC
NC 27713-2852

Tel: +1.919.541.6000
E-mail: rtipress@rti.org
Website: www.rti.org

©2026 RTI International. RTI International is a trade name of Research Triangle Institute. RTI, RTI Press, and the RTI logo are U.S. registered trademarks of Research Triangle Institute. Use of RTI Press publications for text and data mining, AI training, and similar technologies is prohibited without RTI's express written permission.



This work is distributed under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 license (CC BY-NC-ND), a copy of which is available at <https://creativecommons.org/licenses/by-nc-nd/4.0>

<https://doi.org/10.3768/rtipress.2026.op.0100.2603>

www.rti.org/rti-press

Contents

About the Authors	iii
Acknowledgments	iv
Abstract	iv
Introduction	1
Results from Four Approaches to Synthesis	2
Conclusions	8
Data Availability Statement	8
References	9
Appendix A. Acknowledgements of Group Discussion Participants, Manuscript Advisors, and Select Event Staff	9
Appendix B. Process Descriptions and Reflections from Four Approaches to Synthesis	10

About the Authors

Brian Southwell, PhD, distinguished fellow, RTI International. <https://orcid.org/0000-0001-5091-8782> (bsouthwell@rti.org)

Lee Rainie, MA, director, Imagining the Digital Future Center, Elon University

Katie Bowler Young, MFA, senior director, University Collaborations, RTI International

Maria Gallardo-Williams, PhD, director of faculty development and alumni association distinguished undergraduate professor, North Carolina State University

Tony Kipkemboi, BS, staff AI operations engineer, Guild

Jim Redden, BS, director, Technology and AI Automation, RTI International

Matthew Strobl, MBA, MA, founding principal, Trailhead Advisory, LLC

Jon Accarrino, BS, founder, Ordo Digital

Evan Bunnage, BS, software applications programmer, RTI International

Robert Chew, MS, senior research data scientist & fellow, RTI International

Schuyler DeBree, BA, communication strategist, RTI International

Heather Griffiths, PhD, associate professor, Fayetteville State University

Haley Hickman, MLS, executive assistant, Institute for Artificial Intelligence and Emerging Research, North Carolina Central University

Rachel (Ray) Levy, PhD, executive director of Data Science and AI Academy and professor of mathematics, North Carolina State University

Selena Monk, DHSc, instructor, public health education and health administration, North Carolina Central University

Rachel Page, MPH, deputy director, Global NC

Jessica Reif, BS, PhD candidate, Fuqua School of Business, Duke University

Rebecca Ivic, PhD, associate dean of research, scholarship, and creative activity and professor, College of Communication and Information Sciences, University of Alabama

Elizabeth Mannshardt, PhD, vice president and director, Statistics and Data Science Center, Westat

RTI Press Associate Editor

Cynthia Bland

Acknowledgments

The Imagining the Digital Future Center at Elon University and RTI International's Fellows Program and University Collaborations Office organized *The Human Edge: Our Future with Artificial Intelligences*, the convening described in this paper, with support from Burroughs Wellcome Fund, the Knight Foundation, and Schmidt Sciences. We are grateful for comments from two anonymous reviewers whose feedback substantially improved the paper, editing support from Claire Korzen, and graphics support from Nunzio Landi.

Abstract

Although organizations are developing and applying artificial intelligence (AI) technologies in many domains, such as education, health, creative expression, and governance, organizations do not yet collectively agree on useful research questions to guide future inquiry about roles for humans amidst AI. We need empirical evidence comparing different technologies and approaches and assessing both intended effects and unintended consequences. In this paper, we present both substantive findings and methodological lessons learned to illustrate how human and AI-enabled processes together can inform agenda-setting in emerging areas of inquiry. We also spotlight key questions we should be asking about roles for humans amidst societal changes involving AI. We designed and implemented a network-building and agenda-generating effort in fall 2025 to generate key research questions. As part of a hybrid online and in-person event involving nearly 500 participants, 178 in-person attendees participated in one of 23 in-person discussion groups, each moderated by a facilitator with social science research experience and group interviewing skills. Following the discussions, multiple teams of analysts used AI tools and group discussion to summarize key questions from the convening. Across teams, recurring dilemmas, opportunities, and questions emerged, including questions involving comparisons and tradeoffs among different goals for AI use, variations in how people may experience or be affected by AI, and how to optimize human-AI interactions to accomplish tasks. Iterative review and human judgment at multiple stages demonstrated value in all analyses, regardless of the specific AI tools employed, illustrating how human-in-the-loop design can be useful.

Introduction

Despite a flurry of headlines regarding the arrival of technologies involving various forms of artificial intelligence (AI), public discourse has not focused extensively on ethical and logistical questions about human interactions with the range of new entities that likely will emerge in the coming decades, as Boyle (2024) notes in *The Line: AI and the Future of Personhood*. Commentators have celebrated, bemoaned, and speculated about the arrival of AI, and recent public opinion research suggests widespread popular concern about the adoption of AI tools (Rainie & Anderson, 2025). Some editorials and review papers to date have focused on emerging technical issues specific to topical domains without consideration of what kinds of questions could and should be asked about human choices, agency, and opportunities. In other cases, we can find broad philosophical exploration about the nature of humans and technology that lacks specific suggestions for steps that organizational decision-makers can take. We need formalized research agendas in many areas, including education, health, creative expression, and governance.

Generating and analyzing informative and coherent evidence to address concrete questions stemming from public commentators' observations and concerns remains a task for scientists to address. Researchers have begun to ask relevant questions, such as whether we should include AI applications in organizational charts (Burton et al., 2024), how generative AI applications can enhance organizational efficiency (Obrenovic et al. 2025), and the extent to which AI applications can improve patient outcomes in health care (Gala et al., 2024). We need empirical evidence comparing different technologies and approaches and assessing both intended effects and unintended consequences. By presenting both substantive findings and methodological lessons learned here, we illustrate how human and AI-enabled processes together can inform agenda-setting in emerging areas of inquiry, and we spotlight key questions we should be asking about roles for humans amidst societal changes involving AI. Research to generate such evidence will require we continue to

articulate and assess answerable questions about technologies that are not yet altogether here and human capacities and preferences, which also can change over time.

How to Build a Research Agenda Collaboratively

Efforts to address possibilities for human interaction with AI applications can benefit from multiparty collaboration, particularly when researchers seek to establish shared lines of inquiry. Such collaboration offers promising diversity and greater volume of ideas but also introduces challenges. Collaboration is not always necessary for adequate task completion, as individual researchers or organizations can make substantial progress on their own in many instances. Collaboration also can offer an illusion of progress if not aligned with organizational goals and well-organized (Antikainen et al., 2010; Gardner, 2005; Schmutz et al., 2024). Agreement regarding participant motives and goals can help sustain collaborative efforts beyond a meeting or two, ensuring momentum after initial meetings and setting groundwork for concrete outcomes. Additionally, an important aspect of team collaboration performance is the interdependence of tasks and the group's sense of distinctiveness and shared identity (Gundlach et al., 2006). Seeking such group identity among researchers and practitioners can encourage continued efforts to answer questions that matter by facilitating forums for discussion and inquiry.

We convened relevant professionals to generate and refine questions about human interactions with AI tools and to discuss research needs associated with developments in education and workforce training, health care, creative expression professions, and governance. In analyzing the results of the collaborative discussions that ensued, we also explored ways groups of people can work together to use AI tools to make sense of human brainstorming inputs. In other words, we were able to assess some approaches to generating and developing questions about human interaction with AI. What we report here is both a summary of those efforts and reflections on implications for future research.

An Agenda-Setting Convening to Ask Questions about Human Interaction with AI

We designed and implemented a network-building and agenda-generating effort in fall 2025. In September 2025, approximately 450 researchers, practitioners, AI industry experts, leaders from eight universities in North Carolina (Duke University, Elon University, Fayetteville State University, North Carolina A&T State University, North Carolina Central University, North Carolina State University, the University of North Carolina at Chapel Hill, and the University of North Carolina at Greensboro), the National Humanities Center, and a variety of professional organizations participated either online or in-person at a summit organized by RTI International and Elon University. In the afternoon, we invited the 178 in-person attendees to participate in one of 23 in-person discussion groups, each moderated by a facilitator with social science research experience and group interviewing skills. Each group primarily had an opportunity to focus on one of four topic areas: education, workforce, and labor; human agency and creativity; mental and physical health and well-being; and governance and democracy. Participants indicated the topic they hoped to discuss before the convening, and all participants joined a discussion in which they had expressed interest.

Participants discussed “the most important questions we should be asking and addressing about the role of humans in this subject area as AI systems continue to evolve,” guided by moderators who used a discussion guide developed by conference organizers in advance. Moderators also invited participants to discuss who is well-positioned to examine these questions and what resources are needed to pursue answers to key questions. Groups submitted answers using a shared electronic form that recorded text responses in fields in an Excel file. At the day-long event, organizers used AI tools in combination with human discussion (as described in a later section of this article) to generate an initial summary of ideas, which was shared with the entire group of in-person attendees within an hour of completion of the discussions. After the event, organizers invited all participants to

consider joining one of multiple teams of analysts to use tools and group discussion to summarize key questions from the convening. What follows is a description of the analysis efforts conducted by four analysis teams. Reflections from each team on the process of distilling insights using various approaches also appear in Appendix B.

Results from Four Approaches to Synthesis

Four human teams comprising participants from the convening used different tools and processes to generate final sets of research questions based on convening input. All four teams had access to the same file, which contained text inputs from the 23 groups. We present the final set of questions from each of the four analysis teams after a brief description of each unique analytic approach. Table 1 summarizes different analytic approaches. Additional details on analysis differences and team observations appear in Appendix B.

1. Rapid Processing and Follow-up Using SurveySynth and RTI GPT (Jim Redden, Evan Bunnage, and Jessica Reif)

Brief Description of Analysis Approach

The team used a combination of human judgment and AI tools developed for internal use at RTI to synthesize rapidly, giving an initial presentation back to participants within an hour of the discussion exercise concluding. The approach used parallel processing with multiple validation checkpoints. When group discussions concluded on the day of the event, three team members downloaded submitted data and simultaneously initiated three distinct work streams. One researcher employed SurveySynth, an RTI-developed AI tool designed to qualitatively code and synthesize free-text survey responses. A second researcher used RTI GPT, the organization's internal large language model (LLM) tool based on OpenAI's GPT-4o and GPT-4o mini, applying a detailed prompt to synthesize the breakout responses. In parallel, a third researcher conducted a complete human review of the entire group input dataset, reading through all submitted responses without AI assistance.

Table 1. Summary of Analysis Approaches

Team	Key Analysis Attributes
Redden, Bunnage, & Reif	Team provided initial, rapid report on day of the event. Three researchers each used one of three approaches: (1) used RTI's SurveySynth to code and synthesize free-text survey responses, (2) used RTI GPT, a tool based on OpenAI's GPT-4o and GPT-4o mini, to synthesize responses, or (3) conducted human review, reading responses quickly without AI assistance. The team compared results and met to refine results further after the day of the event.
Gallardo-Williams, Chew, & Levy	The team worked together first to develop a criteria rubric grounded in established research design principles and then instructed GPT-5 (model alias:gpt-5-2025-08-07) to score each research question submitted by using the rubric.
Strobl, Griffiths, Page, & Hickman	Team together prompted ChatGPT Plus 5.0 for summary results and then assessed results using virtual meetings to facilitate collaborative results review by team members.
Kipkemboi, Accarrino, DeBree, & Monk	Each team member separately used Python scripts, Claude 3.5 Sonnet, Microsoft Copilot, or a combination of these to analyze one of the topical themes, personally judging data filtering and extraction results from an AI tool to then refine results for group discussion and summary.

Identified Questions: Education, Workforce, and Labor

- What foundational and advanced human skills—such as soft skills, critical thinking, and creativity—should educators and employers prioritize to ensure humans thrive alongside AI in the future workforce?
- How can AI-enabled personalization in education and workforce systems balance innovation and fairness to ensure equitable access and avoid exacerbating societal inequalities?
- What ethical frameworks, regulatory structures, and institutional collaborations are essential to ensure responsible, transparent, and inclusive adoption of AI across education and workforce sectors?

Identified Questions: Mental and Physical Health and Well-Being

- How can we develop AI tools in health care that prioritize transparency, accountability, and trustworthiness while ensuring equitable access for underserved populations?
- What strategies and partnerships are required to train health care professionals, patients, and communities in responsible and effective AI integration, fostering collaboration across diverse expertise and perspectives?

- How can AI systems balance efficiency and human-centric care outcomes while minimizing environmental costs and resisting profit-driven incentives in health care innovation?

Identified Questions: Human Agency and Creativity

- How can we preserve human agency and ensure equitable representation in creative and knowledge fields as AI becomes an active collaborator in cultural and creative pursuits?
- What policies and frameworks are necessary to address intellectual property in AI-human collaborations, protect human creativity, and prevent cultural commodification?
- How do we redefine the relationship between productivity and creativity in a society increasingly shaped by AI, while mitigating environmental and societal harms?

Identified Questions: Governance and Democracy

- What regulatory frameworks at local, national, and global levels can balance the need for AI transparency and accountability with fostering innovation across diverse sectors?

- How can affected communities, civil society organizations, and governing bodies meaningfully participate in AI-related governance processes to ensure equity, oversight, and ethical decision-making?
- What strategies can promote civic knowledge, media literacy, and democratic engagement in the age of increasingly fragmented, AI-influenced information landscapes?

2. A Rubric Approach Using GPT-5 Emphasizing Theme Frequency (Maria Gallardo-Williams, Rob Chew, and Ray Levy)

Brief Description of Analysis Approach

The team's approach involved developing a rubric grounded in established research design principles and then instructing GPT-5 (model alias:gpt-5-2025-08-07) to score each research question generated during the workshop using this rubric. The rubric or framework draws from leading methodological sources, including Creswell and Creswell (2018), Booth et al. (2016), and the FINER (Cummings, Browner, & Hulley, 2013) and SMART (Adeoye and Adong, 2023) evaluation models. These foundations emphasize that a high-quality research question should be clear, researchable, significant, complex, and feasible while maintaining ethical integrity. The team's evaluation system operationalized these qualities into seven measurable criteria: (1) clarity and focus, (2) ability to be researched, (3) significance and relevance, (4) theoretical or practical contribution, (5) complexity and depth, (6) feasibility and scope, and (7) ethical soundness. The team used the LLM to score the questions from the workshop discussions on a 1 to 5 scale for each criterion, and the two highest-scoring questions per topic were selected to represent each domain. Human analysts were responsible for developing the rubric and prompts, and the LLM collated, scored, and evaluated the research questions.

A distinctive feature of the team's approach is the inclusion of a frequency dimension as a proxy for expressed participant interest and perceived importance. Because workshop participants often

generated overlapping or near-duplicate questions, frequency counts captured how often similar ideas emerged across breakout groups. The team also used frequency as a tiebreaker consideration when multiple questions received identical high scores.

Identified Questions: Education, Workforce, and Labor

- When and how should humans remain “in the loop” to detect and correct AI errors in education and work?
- What strategies ensure equity, infrastructure, and access for all communities in AI-enabled education?

Identified Questions: Mental and Physical Health and Well-Being

- What AI prompt best practices are most impactful for reducing bias in health research?
- How can we ensure data protection and privacy and assess the trustworthiness of AI systems?

Identified Questions: Human Agency and Creativity

- How do we avoid deepening inequalities caused by AI?
- Who holds creative license or intellectual property rights over AI-generated work—the AI, the prompter, or others?

Identified Questions: Governance and Democracy

- How do we ensure that automated decision-making is fair, equitable, and just and that it advances human flourishing?
- How do we implement AI governance that ensures a human-centered, ethical focus at every stage—idea, design, development, and deployment?

3. Iterative Development of Theses and Questions Aided by ChatGPT Summary (Matthew Strobl, Heather Griffiths, Rachel Page, and Haley Hickman)

Brief Description of Analysis Approach

To analyze the group input data, the team used ChatGPT Plus 5.0 in conjunction with Google Meet to facilitate collaborative discussion and review among team members. Human analysts played an active role throughout the process, guiding the prompts, refining inputs, and interpreting the outputs. The team instructed the tool not to offer praise for our work in a subsequent prompt to test tone and framing. When ChatGPT's responses stalled or failed to complete, the team removed or simplified the more-complex requests to maintain progress. At times, when the tool prompted for a response to continue or update its analysis, they intentionally paused to refine instructions and approach. Human feedback occurred organically through team dialogue and is therefore difficult to fully capture. They also re-prompted ChatGPT to take on a skeptic's role to challenge and assess its own outputs more critically. When this approach caused repeated stalling, they adjusted the prompt to position it as a neutral judge instead, offering more-detailed guidance on how to analyze the data. Ultimately, human analysts exercised judgment at each stage to determine which iterations produced the most-accurate, most-balanced, and most-insightful analysis.

Identified Theses and Questions: Education, Workforce, and Labor

Thesis 1: The degree of alignment between educational curricula and employer-defined AI competencies predicts the distribution of re-skilling opportunities.

- *Relevant questions:* How closely do education and training programs map to employer-identified AI skill needs? What groups or regions show the largest gaps between available training and workforce demand? What mechanisms (curriculum audits, employer feedback loops) improve skill alignment?

Thesis 2: The presence of structured partnerships among education systems, employers, and public agencies influences implementation of AI-related training programs.

- *Relevant questions:* What forms of cross-sector partnerships are associated with faster adoption of AI training modules? How do funding, governance, or shared infrastructure arrangements affect implementation scope? What indicators show successful program uptake across sectors?

Identified Theses and Questions: Mental and Physical Health and Well-Being

Thesis 1: The adoption of AI-supported triage or scheduling tools in community health centers changes referral completion and care coordination patterns.

- *Relevant questions:* How do AI scheduling or triage tools affect patient referral completion rates and wait times? What differences appear between centers with and without AI-enabled coordination? How do staff workflows or communication practices change after implementation?

Thesis 2: Governance transparency in AI-driven health platforms affects data-sharing frequency between institutions.

- *Relevant questions:* Which governance practices (e.g., consent visibility, audit reporting) are linked to higher data-sharing frequency? How do institutions evaluate trade-offs between transparency, privacy, and efficiency? What evidence links governance changes to shifts in data exchange activity?

Identified Theses and Questions: Human Agency and Creativity

Thesis 1: The level of human review required in AI-assisted content production affects decision turnaround and attribution practices.

- *Relevant questions:* How does the number or type of human review steps correlate with output speed and approval rates? What documentation practices

assign authorship in AI-assisted outputs? How do patterns differ across organizational or sectoral contexts?

Thesis 2: Institutional rules defining AI authorship determine the approval chain and accountability structure in creative outputs.

- *Relevant questions:* What policies outline when AI-generated material requires human sign-off? How do institutions document accountability for AI-assisted decisions? What variations exist in governance across industries using AI in creative or analytical work?

Identified Theses and Questions: Governance and Democracy

Thesis 1: The proportion of civil society representation in AI oversight bodies correlates with the perceived fairness of governance outcomes.

- *Relevant questions:* What percentage of governance seats are occupied by civil society actors? How do participants assess fairness and responsiveness of oversight decisions? Which participatory mechanisms align with perceptions of procedural fairness?

Thesis 2: The communication framing of AI regulation affects citizen perceptions of institutional legitimacy.

- *Identified questions:* How do alternative framings (risk-based, rights-based, innovation-based) influence perceived legitimacy? What messaging strategies improve understanding of regulatory intent and accountability? How do demographic or ideological factors mediate public responses to AI policy communication?

4. Human-Guided Use of Python Scripts, Claude 3.5 Sonnet, and Microsoft Copilot to Analyze Data Filtering and Extraction (Tony Kipkemboi, Jon Accarrino, Schuyler DeBree, and Selena Monk)

Brief Description of Analysis Approach

Each of four analysts took responsibility for one of the topical themes and selected a strategy for analysis.

To analyze the data for the Education, Workforce, and Labor track, an analyst used Copilot and Excel

in combination with a structured prompt-based approach, first uploading the breakout session notes captured in the Excel file into Copilot and asking for extraction and filtering of all responses tagged under the relevant topic area. Copilot helped identify and synthesize key themes across multiple discussion summaries when asked to generate research questions. The analyst then determined which insights were most representative of the group's concerns.

To analyze the Mental and Physical Health and Well-Being track, one analyst used a systematic iterative approach using Anthropic's Claude, an AI assistant, combined with manual review of the data, first extracting all group responses from this topic area and asking Claude to identify recurring themes. The analyst iteratively refined analysis by asking follow-up questions about specific tensions, such as, "How do the discussions characterize the tension between profit motivation and patient welfare?"

To analyze the data for the Human Agency and Creativity track, one analyst combined Claude with Python-based data analysis, first uploading the Excel file containing breakout session notes and using Python scripts to filter and extract all four responses tagged under "Human agency and creativity." The analyst prompted Claude to conduct a comprehensive thematic analysis and then reviewed the thematic synthesis and made critical judgments about which themes were most prevalent and pressing across discussions.

To analyze the Governance and Democracy track, one analyst used Claude (Anthropic, 2025) in combination with a structured comparative framework, isolating the group governance-focused responses and prompting Claude to map the landscape of discussion, identify philosophical positions, and highlight tensions. The human analyst then assessed results and asked Claude to compare different governance philosophies evident in the discussions (e.g., "Should there be standards at all?" vs. calls for comprehensive regulation).

Identified Questions: Education, Workforce, and Labor

- How can education systems and workforce development programs evolve to prioritize and preserve uniquely human skills, such as critical thinking, empathy, adaptability, and creativity, in an AI-augmented world?
- What policies, partnerships, and infrastructures are needed to ensure equitable access to AI tools and education while preventing the deepening of digital divides and ensuring that all communities benefit from AI advancements?

Identified Questions: Mental and Physical Health and Well-Being

- How can health care systems design and implement AI tools that preserve clinician autonomy, critical thinking capacity, and patient-centered care while addressing the tension between profit-driven efficiency metrics and humanistic health outcomes, particularly in ensuring that underrepresented populations and low-resource contexts are meaningfully included in data, design, and deployment decisions?
- What regulatory frameworks, transparency standards, and collaborative governance structures are needed to ensure AI-augmented health care maintains privacy protections, addresses environmental costs, balances the promise of synthetic data with the need for prospectively collected real-world evidence, and establishes clear accountability when AI-driven decisions lead to patient harm?

Identified Questions: Human Agency and Creativity

- As AI systems become increasingly integrated into creative and knowledge work, what foundational human capabilities and "starter skills" must be preserved and cultivated to maintain meaningful agency over decision-making and creative processes, and how can we design human-AI partnerships that enhance rather than erode human cognitive autonomy and creative capacity?

- What governance frameworks, attribution systems, and equity-centered policies are necessary to prevent AI from exacerbating marginalization, commodifying culture, and concentrating creative agency in corporate entities, while ensuring that diverse communities maintain ownership, voice, and fair compensation in the development and outputs of AI-augmented creative work?

Identified Questions: Governance and Democracy

- What multilevel governance frameworks (local, national, international) are necessary to ensure that AI systems deployed in democratic societies maintain transparency, enable meaningful human oversight of automated decision-making, and build public trust while balancing the rapid pace of technological innovation with the need for enforceable standards that protect human rights and prevent the concentration of power in corporate entities?
- How can democratic societies cultivate widespread AI literacy and civic knowledge necessary for informed participation in AI governance, ensure that affected communities have meaningful voice in algorithmic decision-making that affects their lives, and preserve the capacity for critical engagement with information in an environment increasingly saturated with AI-generated content and misinformation?

Across all four topical areas, several cross-cutting themes emerged that reflect shared concerns about the future of human–AI interaction. Participants emphasized the need for transparency, equitable access, meaningful human oversight, and reliable governance structures, regardless of domain. Concerns related to infrastructure and the preservation of human agency also emerged. Highlighting these patterns helps place these domain-specific questions within a broader research landscape involving AI and demonstrates the contribution of coordinated inquiry.

Conclusions

A collaborative effort to bring together experts, researchers, and practitioners to identify useful research questions about roles for humans working in various domains amidst AI systems generated a robust research agenda. Nearly 200 in-person participants working in 23 groups contributed sets of questions to guide future research focused on four different topical areas. Small teams of analysts then used a combination of human review and AI-supported tools to organize and assess these inputs, identifying key questions to inform an agenda for future inquiry. That process yielded relevant questions that could offer a foundation for future studies and collaborative projects.

Across teams, recurring dilemmas, opportunities, and questions emerged. These included questions involving comparisons and tradeoffs among different goals for AI use, variations in how people may experience or be affected by AI, and how to optimize human-AI interactions to accomplish tasks. At the same time, different analytic approaches also yielded a diverse array of potential questions reflecting a range of different evaluation criteria, such as the frequency with which ideas appeared across participant groups or the extent to which a topic lent itself to empirical evidence. Across methods, iterative review and human judgment at multiple stages demonstrated value, regardless of the specific tool employed, illustrating how human-in-the-loop design can be useful. That observation has implications for how we could work optimally with AI tools. Explicitly designing and articulating decision-making processes and assigning humans to not only review output but also initiate and consider the questions that guide projects will be important. Planning workflow processes involves choices about how to involve humans. In this way, human actors can integrate AI tools into work while staying in command of key decision-making, as scholars have recently recommended (e.g., “Artificial Intelligence Is Not a Substitute,” 2025).

Organizations involved in education and labor, health care and well-being, creative expression, and governance may identify opportunities to generate

evidence to address questions such as those posed here. Rather than solely relying on anecdotal observations about single user experiences with available tools, researchers and practitioners may better inform future decision-making about AI engagement by contributing to and learning from literature that addresses such problems. Learning networks of collaborators, including human teams using AI-assisted tools, could advance available literature but will do so most productively in carefully organized partnerships when party motives are clear and task goals are explicit. Without collaborative efforts to address key research questions about human interactions with AI, organizational missteps, inefficiencies, and ethically problematic consequences are possible. With collective efforts to address key questions, leaders of organizations employing humans should be able to make better choices about when and how to engage AI.

Data Availability Statement

Data availability statement: De-identified data supporting the current study are available from the authors upon reasonable request.

References

- Adeoye, M. A., & Adong, C. (2023). The power of precision: Why your research focus should be SMART? *Journal of Education Action Research*, 7(4), 569–577. <https://doi.org/10.23887/jear.v7i4.69757>
- Antikainen, M., Mäkipää, M., & Ahonen, M. (2010). Motivating and supporting collaboration in open innovation. *European Journal of Innovation Management*, 13(1), 100–119. <https://doi.org/10.1108/14601061011013258>
- Artificial intelligence is not a substitute for human intelligence [Editorial]. (2025). *Nature Reviews Psychology*, 4, 753–754. <https://doi.org/10.1038/s44159-025-00517-y>
- Booth, A. (2016). Searching for qualitative research for inclusion in systematic reviews: a structured methodological review. *Systematic Reviews*, 5, 74. <https://doi.org/10.1186/s13643-016-0249-x>
- Boyle, J. (2024). *The Line: AI and the Future of Personhood*. MIT Press.
- Burton, R.M., Obel, B., & Håkonsson, D.D. (2024). Expanding the organizational design space: the emergence of AI robot bosses. *Journal of Organizational Design*, 13, 13–22. <https://doi.org/10.1007/s41469-023-00155-9>
- Creswell, J. W., & Creswell, J. D. (2017). *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*. Sage Publications.
- Cummings, S. R., Browner, W. S., & Hulley, S. B. (2013). Conceiving the research question and developing the study plan. In S. B. Hulley, S. R. Cummings, W. S. Browner, D. G. Grady, & T. B. Newman (Eds.), *Designing Clinical Research* (4th ed., pp. 14–22). Wolters Kluwer/Lippincott Williams & Wilkins.
- Gala, D., Behl, H., Shah, M., & Makaryus, A. N. (2024). The role of artificial intelligence in improving patient outcomes and future of healthcare delivery in cardiology: A narrative review of the literature. *Healthcare*, 12(4), 481. <https://doi.org/10.3390/healthcare12040481>
- Gardner, D. (2005). Ten lessons in collaboration. *OJIN: Online Journal of Issues in Nursing*, 10(1), 1. Available at <https://ojin.nursingworld.org/table-of-contents/volume-10-2005/number-1-january-2005/ten-lessons-in-collaboration/>
- Gundlach, M., Zivnuska, S., & Stoner, J. (2006). Understanding the relationship between individualism-collectivism and team performance through an integration of social identity theory and the social relations model. *Human Relations*, 59(12), 1603–1632. <https://doi.org/10.1177/0018726706073193>
- Obrenovic, B., Gu, X., Wang, G., Godinic, D., & Jakhongirov, I. (2025). Generative AI and human-robot interaction: Implications and future agenda for business, society and ethics. *AI & Society*, 40, 677–690. <https://doi.org/10.1007/s00146-024-01889-0>
- Rainie, L., & Anderson, J. (2025). *Many Americans expect AI to have significant negative impact on human capacities and behaviors such as social and emotional intelligence, analytical thinking and agency by 2035*. Elon University's Imagining the Digital Future Center. Available at <https://imaginingthefuture.org/wp-content/uploads/2025/08/LLM-survey-public-views-on-being-human-in-2035-9-17-25.pdf>
- Schmutz, J. B., Outland, N., Kerstan, S., Georganta, E., & Ulfert, A.-S. (2024). AI-teaming: Redefining collaboration in the digital era. *Current Opinion in Psychology*, 58, 101837. <https://doi.org/10.1016/j.copsyc.2024.101837>

Appendix A. Acknowledgements of Group Discussion Participants, Manuscript Advisors, and Select Event Staff

Jon Accarrino, **Haya Ajjan**, **Mustafa Akben**, **Jeffery Alexander**, Jiyeon An, Daniel Anderson, Valerie Anderson, Warwick Arden, Elena Arthur, **Scott Babwah Brennen**, Ed Balleisen, Nicole Barnes, Tiffany Barnes, Robert Beach, **Brinnae Bent**, Cam Bernstein, Sambit Bhattacharya, William Gene

Black, Stacye Blount, Dillan Bno-Lunn, Georgiy Bobashev, Ganesh Bora, Kristi Boyd, Kevin Boyd, Kerry Branon, Kerry Brock, Marquita Brogdon, Matthew Brophy, **Ramsay Brown**, Ricardo Budet, Yuxuan Cai, Catherine Chao, Katelyn Chedraoui, **Helen Chen**, **Rob Chew**, Kristen Christman, Brianna

Clarke-Schwelm, Alisa Clifford, Bill Coffey, Andrew Cox, Cobey Culton, Cristiane Damasceno, Christin Daniels, Tom Darden, Abhik Das, Connor Davis, **Dilsey Davis**, **April Dawson**, L. Shawn DeBerry Johnson, **Schuyler DeBree**, Anthony Dellinger, Joy Dismukes, Su Dong, Layla Dowdy, Robert Duvall, **Shellery Ebron**, Joe Edney, **Bucky Fairfax**, Kristy Farewell, Brian Feeley, Tim Fennell, Samuel Field, Laura Fiorilli-Crews, Steve Folger, Brent Fox, Sherell Fuller, Tim Gabel, Maria Gallardo-Williams, Genevieve Garland, Markeith Gentry, Debranetta Gethers, Sharon Ghazarian, Leanna Giles, Priyanka Ginwal, Anne Glauber, Penny Gordon-Larsen, Siobahn Grady, Frank Graff, Eric Grebing, Jennifer Greer, Heather Griffiths, Eric Hall, Shane Hamstra, Dana Hancock, John Hardin, Matt Hartman, Stephanie Hawkins, Richard Chris Hayes, Meredith Henderson, Haley Hickman, Melissa Hodge-Penn, Nicole Hodges, Erich Huang, James Hunter, **Jason Husser**, Rebecca Ivic, Daniel Jacobson, **Mohammad Hossein Jarrahi**, Bridget Johnson, Drew Johnson, Eric Johnson, Eric Otto Johnson, Zakaria Jouaibi, Annice Kim, **Tony Kipkemboi**, Alex Kral, Shannon Krauss, Christopher Krebs, Rashmi Kumari, Mark Kurt, Kuldip Kuwahara, Barrot Lambdin, Leo Lambert, Rachel (Ray) Levy, Mark Light, Amy Linder, Kun Lu, **Molly Lynch**, Laurell Malone,

Elizabeth Mannshardt, **Jonathan McCall**, **Heather McCullough**, Liz McLaughlin, Reagan McRae, Selena Monk, Krista Moore, Shona Morgan, Brenna Muldavin, **Amanda Nelson**, Patrick Noltemeyer, Jason Norman, Laura Nyblade, Sherine Obare, **Andreas Oranje**, Grier Page, Rachel Page, Michael Pencina, Shanetta Pendleton, Rekha Perumallu, Mary Elizabeth Piecham, Jon Poehlman, Anuja Purohit, Lee Rainie, Jessica Reif, Jeri Roper Miller, Chris Roush, Amy Roussel, Dina Rousset, **Doug Rupert**, **Lisette M. Saavedra**, Nicole Samuel, Allan Sandoval, Blake Schaeffer, Natalie Schulte, Lauren Scroggs, Christy Shaffer, Scott Singeisen, Melissa Skillings, Luke Smith, Mary Southern, Brian Southwell, Tommy Sowers, Kami Spangenberg, Elizabeth Standafer, Rebecca Stanley, Kathleen Stansberry, Amanda Stockwell, Carmen Strigel, Matthew Strobl, Tamara Terry, Bina Thakkar, TJ Tipper, Jacob Travere, Elizabeth von Briesen, Zaida Walker, **Sarah Egan Warren**, Christopher Waters, Paul Weisenfeld, Anna Wetterberg, Tom White, Michael Willoughby, Emily Witthohn, Andrea Woodson-Smith, Dean Woodward, Katie Bowler Young, Ami Zach, Gary Zarkin, Shannon Zenner, Tianduo Zhang

Note: Moderators and co-moderators are denoted with bold text. Ramsay Brown also offered advice on this essay.

Appendix B. Process Descriptions and Reflections from Four Approaches to Synthesis

Rapid Processing and Follow-up Using SurveySynth and RTI GPT

Jim Redden, Evan Bunnage, Jessica Reif

As described in the main discussion, the team's approach centered on parallel processing with multiple validation checkpoints (Figure B1).

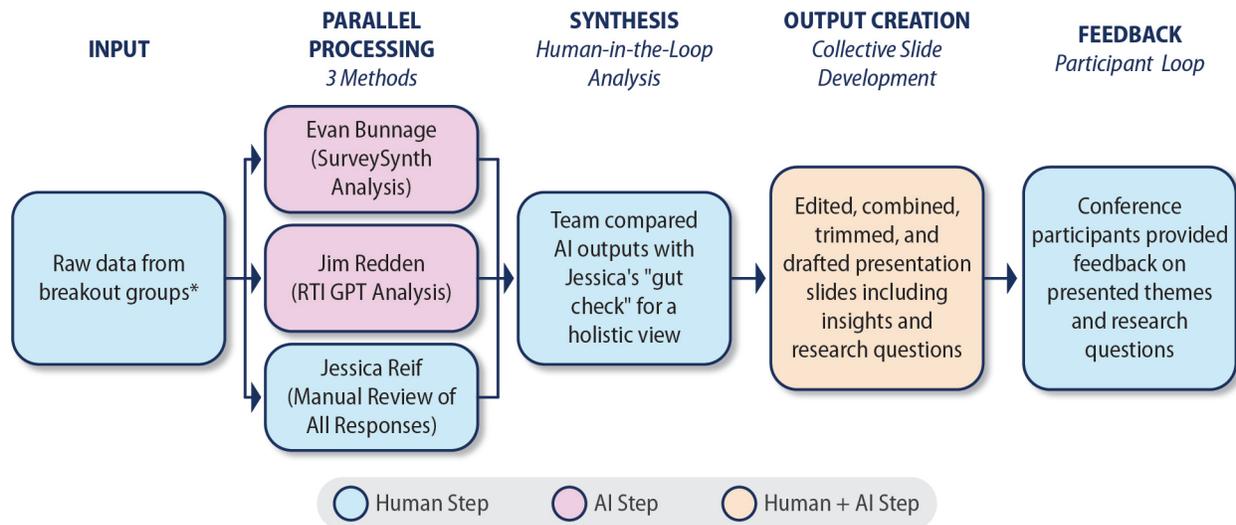
Reflections on the Process

The 30-minute time constraint fundamentally shaped our approach and its limitations. This severe time pressure necessitated human judgment calls that made the process less than fully replicable. We could not systematically document every decision point, curation choice, or selection criterion used to move from raw data to final output. Additionally, the probabilistic nature of large language models (LLMs) means that

identical inputs can produce varying outputs, further complicating reproducibility. These factors underscore an important principle: AI synthesis approaches must be fit for purpose and adapted to specific constraints, objectives, and rigor requirements.

The team's primary objective was to accomplish something genuinely impossible for humans alone within the timeframe: reading, digesting, synthesizing, and creating four production-ready presentation slides from 20 breakout groups in 30 minutes. For rigorous, repeatable scientific qualitative analysis, this would not represent best practice.

However, it proved highly effective for enabling real-time participant engagement with cross-group synthesis, something that added immediate value to the conference experience.

Figure B1. Overview of Plan for Same-Day Synthesis (results in approximately 30 minutes)

*Because three groups submitted after the deadline, the live analysis did not include the full set of breakout teams, which likely led to some of the in-room feedback on the outputs about incompleteness.

One significant limitation emerged from the time constraint itself. Because we began processing at exactly 3:00 p.m. EST during the day of the event, three breakout groups that submitted their synthesis forms late were inadvertently excluded from our analysis. We did not notice the incomplete dataset at the time, which likely explains why some participants reported that the synthesis failed to capture their discussions. This experience highlights how real-time analysis pressures can introduce systematic data collection errors that compromise output quality.

The most valuable insight from this exercise concerns the power of human-AI collaboration over full automation. Maintaining human involvement at multiple stages proved essential. The researcher who read the complete dataset could validate whether AI-generated themes accurately reflected the content. Human researchers could instruct the language model on formatting constraints and make nuanced judgments about which insights to prioritize when condensing five potential themes down to two or three for presentation. These curation decisions required contextual understanding and stakeholder awareness that current AI systems cannot replicate.

This case study demonstrates that AI tools can enable qualitative synthesis tasks that would otherwise be

infeasible because of time or resource constraints. However, the appropriate use of AI for synthesis depends entirely on analytical objectives and quality requirements. For contexts requiring rigorous, systematic, and fully reproducible analysis, removing time constraints and developing more-structured protocols would be advisable. For contexts where rapid synthesis enables valuable real-time engagement and collaboration, as in our conference setting, a hybrid approach with strong human oversight and validation can deliver meaningful results. The key to success lies in designing approaches that strategically combine human and AI capabilities, leveraging each system's strengths while remaining clear-eyed about the limitations and tradeoffs inherent in the chosen methodology.

A Rubric Approach Using GPT-5 Emphasizing Theme Frequency

Maria Gallardo-Williams, Rob Chew, and Ray Levy

The team's approach to evaluating research questions involved developing a rubric grounded in established research design principles and then instructing GPT-5 to score each research question generated during the workshop using this rubric, as described in the main discussion. Human analysts were responsible

for developing the rubric and prompts, and the LLM collated, scored, and evaluated the research questions. A distinctive feature of the approach is the inclusion of a frequency dimension as a proxy for participant interest and perceived importance.

Reflections on the Process

Overall, this structured approach provided a transparent, repeatable, and data-informed method for evaluating research questions that can be efficiently implemented within a workshop setting. Future iterations could benefit from a more-detailed codebook to contextualize each dimension and from human validation of the LLM's scoring. Although principled, the method reflects a pragmatic bias toward utility and measurability, potentially reducing inquiry to what is easily operationalized and privileging questions that align with standard forms of evidence. Pairing this system with mechanisms that intentionally reward novelty and intellectual dissent could yield a richer and more diverse set of final research questions.

Iterative Development of Theses and Questions Aided by Chat GPT Summary

Matthew Strobl, Heather Griffiths, Rachel Page, and Haley Hickman

As described in the main discussion, the team used ChatGPT Plus 5.0 in conjunction with Google Meet to facilitate collaborative discussion and review among team members. Human analysts played an active role throughout the process, guiding the prompts, refining inputs, and interpreting the outputs.

Reflections on the Process

This exercise revealed several insights about how to optimize similar analytical processes in the future. One key challenge was the lack of consistency across notetakers and in the application of definitions, which made it more difficult to align findings and ensure comparability across analyses. Establishing shared frameworks and clearer documentation standards at the outset would help mitigate these inconsistencies. At the same time, the process demonstrated the importance of continuous integration of outputs, reminding us that, in many ways, “the process is the point.” The iterative exchange between human

analysts and the tool added essential depth and adaptability, underscoring that human judgment and synthesis remain critical components of AI-assisted analysis. Finally, our approach reflected a balance between inductive and deductive reasoning. First, we provided broad instructions. Upon reviewing the research questions that ChatGPT generated, we repeatedly refined our requests, which led to ChatGPT stalling. When our prompts caused ChatGPT to stop responding, we rephrased slightly and continued along a different angle. Second, upon successfully producing usable output, we provided triangulating guidance, switching from a deductive approach to a more inductive approach. For example, in one sequence, we requested output in the form of thesis statements. When we were satisfied with the thesis statements, we directed ChatGPT to produce specific, actionable research questions that could resolve the thesis statements, producing research questions that improved upon the earlier iteration.

This progression suggests that flexibility and reflection are vital for producing thoughtful, credible results in such exercises. It also suggests that blunt force questioning produces inferior results relative to our later technique.

Human-Guided Use of Python Scripts, Claude 3.5 Sonnet, and Microsoft Copilot to Analyze Data Filtering and Extraction

Tony Kipkemboi, Jon Accarrino, Schuyler DeBree, and Selena Monk

Four team members used distinct processes for assessment, as described in the main discussion. Analysts used Python, Claude, Microsoft Copilot, or a combination of these, with prompts such as, “Analyze these Human Agency and Creativity responses and identify the major themes, key tensions, stakeholder perspectives, and resource needs across all four discussion groups.” This iterative and collaborative process balancing human insight and AI assistance allowed for both speed and depth in analysis.

Reflections on the Process

For one analyst, the process highlighted the critical importance of transparency and human judgment in AI-assisted qualitative analysis. Unlike closed

systems, using Claude with explicit Python scripts allowed them to see exactly how data was filtered and extracted, which built trust in the analytical process. The AI excelled at identifying patterns and synthesizing themes across multiple texts, a task that would have been time-consuming manually, but the quality of output depended heavily on the analyst's ability to craft effective prompts and critically evaluate results. They found that breaking the analysis into stages (extraction, thematic coding, synthesis, question formulation) prevented the AI from imposing predetermined frameworks and allowed emergent themes to surface authentically from the data. A key insight is that AI tools are most valuable when they make the analytical process more transparent rather than more opaque. The collaborative approach, where AI handles pattern recognition and the human analyst provides domain expertise, contextual interpretation, and final judgment, produced richer results than either human-only or AI-only analysis would have achieved. For future exercises, the analyst recommends that participants be trained not just in using AI tools, but in the critical thinking skills needed to effectively prompt, evaluate, and refine AI outputs, ensuring that technology augments rather than replaces human insight in research synthesis.

For another analyst, the experience using Copilot in Excel highlighted how AI can be a powerful tool in collaborative analysis when integrated into a familiar workflow and tailored to the user (Microsoft, 2025). One analyst has used Copilot extensively and noted that Copilot has adapted to their style and communication preferences, allowing it to generate outputs that reflect their voice and thinking process. This process made the work efficient and authentic, enabling the analyst to focus on higher-order judgment and interpretation rather than manual sorting or summarizing. In contrast, tools used by other teams, such as the RTI-adapted tool described elsewhere that had not yet been optimized for use by this analyst at the time of their work, appeared to produce outputs that missed key themes deemed important by this team and results that required significant reworking, which added to the workload and raised important questions about AI-assisted analysis. This comparison suggested that AI tools

are most effective when they are embedded in our workflow, learn from iterative use, and support, not replace, the human insight. It also suggests that for future exercises, we should prioritize tools that allow for customization, transparency, and collaboration, rather than one-size-fits-all automation.

Large Language Models Used in Analysis

Anthropic. (2025). Claude 3.5 Sonnet.

<https://www.anthropic.com/claude>

Microsoft. (2025). Copilot (GPT-5).

<https://copilot.microsoft.com/>

OpenAI. (2025). GPT-4o. <https://platform.openai.com/docs/models/gpt-4o>

OpenAI. (2025). GPT-4o mini. <https://platform.openai.com/docs/models/gpt-4o-mini>

RTI International is an independent scientific research institute dedicated to improving the human condition. We combine scientific rigor and technical expertise in social and laboratory sciences, engineering, and international development to deliver solutions to the critical needs of clients worldwide.

www.rti.org/rtipress

RTI Press publication No. OP-0100-2603