

# **AI-centered Communications Consulting with Civil Society Organizations: Case Studies and Preliminary Lessons Learned**

John Wihbey, Caleb Okereke, and Yash Phalle<sup>1</sup>

Working Paper

Last updated: July 24, 2025

AIMES Lab  
Northeastern University



---

<sup>1</sup> John Wihbey, Associate Professor [[j.wihbey@northeastern.edu](mailto:j.wihbey@northeastern.edu)]; Caleb Okereke, Doctoral Researcher; Yash Phalle, Master's Candidate in Artificial Intelligence

Table of Contents

Executive Summary..... 2

Introduction.....3

Contextualization and Literature Review..... 3

Motivating Questions.....4

Case Studies..... 5

Lessons Learned.....6

Recommendations..... 9

Future Questions and Directions..... 11

References.....13

## Executive Summary

This paper examines how generative AI tools can support civil society organizations through a study of academic consulting engagements with five Boston-area nonprofits from January to April 2025. These academic consulting sprints, conducted directly by 15 Northeastern University graduate students in a course setting and the authors of this report, were focused on communications and media work and used an action research approach. A central finding is the concept of "organizational legibility to AI"—the challenge that civil society organizations' unique contexts, values, and relationship-based communication strategies are not immediately interpretable by AI systems, requiring significant human mediation to make organizational realities comprehensible to these tools. Key takeaways include that AI tools became more effective only after consulting teams gained extensive domain knowledge about each organization, that stakeholders sometimes overestimated AI capabilities, and that hybrid approaches combining traditional methods with AI assistance worked best for community-facing work.

The paper's recommendations emphasize investing in AI literacy before deployment, starting with lower-risk documentation tasks, maintaining human oversight for mission-critical communications, developing workflows that preserve authentic engagement while leveraging AI efficiency, and focusing on sustainable integration approaches that do not exceed existing organizational capacity. We conclude that successful AI adoption in civil society contexts requires treating AI as a collaborator rather than replacement, enhancing rather than automating the relational work that defines effective nonprofit communication.

## Introduction

The promise of generative AI tools to help both large business enterprises and commercial startups has dominated public discourse in recent years, but the utility of these new technologies for smaller, more boutique nonprofits and civil society organizations remains unclear. Issue-specific organizations in the public and social sector have narrowly defined or even idiosyncratic requirements, often operate in local contexts, and typically have highly specific stakeholder groups and intended audiences. They often lack technological capacity as the firms themselves are small and have low head count.

From January to April 2025, a group of 15 graduate students at Northeastern University, in partnership with the AI-Media Strategies Lab (AIMES), performed service project-related communications consulting for five civil society organizations in the metro Boston area in order to explore the utility of generative AI tools. The substantive work in these engagements effectively lasted 12 weeks and therefore constituted “sprints” or interventions of short duration. As part of the the “AI and Media Industries” graduate course and Northeastern’s Service-Learning program (part of Community-Engaged Teaching and Research)<sup>2</sup>, teams of three students worked with each external partner client to assess needs, identify workflows where AI technologies might be helpful, and consult with personnel in the civil society organizations to help them with diverse communications- and media-related needs. For a framework for their engagement, the academic consulting teams operated using an action research approach, which involves general cycles of four phases – plan-act-observe-reflect (Stringer & Aragón, 2021; Coghlan, 2011).

This paper seeks to highlight lessons and qualitative observations from these exploratory consultative partnerships. Each of the organizations has been anonymized here in this discussion, although we provide a sense of the domain in which each civil society group is operating. We seek to share learnings with the wider fields of communication and media, where this new generation of technologies are prompting many to ask a converging set of questions that extend across workstreams in news media, social media, advertising, public relations, and strategic communication (Guzman & Lewis, 2024).

## Contextualization and Literature Review

The recent surge in generative AI tools, including large language models (LLMs) like GPT-4 and Claude 4, has opened new possibilities across the communication landscape. Scholars have long anticipated a range of new communications AI-powered capabilities, from automated content production to strategic messaging and audience engagement (Bommasani et al., 2021; Kietzmann et al., 2018). While large corporations and media firms have rapidly begun integrating these tools into their core functions, much of the discourse centers around commercial scalability and high-resource contexts (O’Reilly Media, 2022; Zhang et al, 2023). By contrast, there remains a gap

---

<sup>2</sup> <https://cetr.northeastern.edu/>

in understanding how such tools may serve smaller, mission-driven organizations with constrained budgets, limited technical expertise, and context-specific communication needs.

Civil society organizations and nonprofits often approach technology through a lens of pragmatism, rather than innovation for innovation's sake. Research on technology adoption in nonprofits emphasizes how limited staff capacity, low digital literacy, and mission-centered rather than product-centered priorities shape their engagement with new tools (Burt & Taylor, 2003; Gálvez-Rodríguez, Caba-Pérez, & López-Godoy, 2016). Many rely on ad hoc strategies for communication and storytelling, privileging authenticity and community trust over automation and reach (Lovejoy & Saxton, 2012). This presents both a challenge and an opportunity: generative AI could streamline operations, yet risks eroding the contextual nuance and relational ethos that define these organizations' public engagement.

Moreover, there is a growing body of work that examines participatory and consultative models for tech support in nonprofit contexts. Approaches like service-learning and co-design emphasize mutual learning, knowledge exchange, and power-sharing between technically skilled actors (e.g., students or researchers) and civil society stakeholders (DiSalvo et al., 2013; Light & Miskelly, 2015). These frameworks are especially salient when working with generative AI tools, whose outputs are shaped not only by data but by human prompts, use patterns, and editorial values. They also foreground the importance of process over product; that is, of building understanding and workflows that extend beyond delivering technological solutions.

Finally, studies of workflow and organizational integration point to a “last mile” problem in AI implementation: even when tools are technically useful, they often fail to align with how small teams operate day-to-day (Hussain & Rizwan, 2024; West, 2018). Nonprofits and issue-specific civil society groups face especially high barriers to sustained adoption, including unclear ROI, lack of IT support, and ethical considerations around surveillance, representation, and voice (Eubanks, 2018). This paper contributes to that body of work by offering qualitative insights from consulting relationships that foreground these challenges and opportunities in real-world contexts.

## Motivating Questions

The diverse set of experiments embodied in the various service-learning engagements between graduate students and civil society organizations did not directly test hypotheses or engage in sustained data collection and field research. Rather, the engagements were action research-based, and they were meant to be practical in application and qualitative in their takeaways. The engagements were designed by the researchers and guided by the following general questions:

- To what extent do generative AI tools align with the values, priorities, and constraints of issue-specific civil society groups?

- How does generative AI affect the relational nature of communication work in nonprofits, where trust, authenticity, and local context are paramount?
- How might the challenges faced by small civil society organizations reflect or differ from those of larger institutions in the adoption of generative AI?
- What types of tasks or workflows are most effectively supported by generative AI in mission-driven nonprofit contexts?

Again, these questions helped to frame the approach, but groups did not systematically collect data as they would in a more formal research engagement.

## Case Studies

The five different engagements with civil society organizations took diverse shapes, and many required significant iterative steps in the beginning to establish alignment and specify desired deliverables. The following constitute the five case studies in their broad outlines:

### **Organization 1 | *Local Public Agency***

An academic consulting team collaborated with a municipal housing agency to strengthen its digital literacy program for elderly residents. The team supported both wellness and afterschool workforce initiatives by developing a sponsor list with relevant contacts and partnership details. Their primary focus, however, was on enhancing a 7-week computer skills course for seniors. Using an action research approach, the team co-created a low-tech, user-friendly support website that consolidated key tools and lessons from the course, designed to serve as a long-term resource for learners. Leveraging AI tools like ChatGPT and Claude, they analyzed existing materials, filled content gaps, and produced printable guides and tiered activities tailored to varying skill levels.

### **Organization 2 | *Youth-focused Arts Organization***

The academic consulting team partnered with a community-based youth arts organization focused on empowering young people through storytelling and performance. The team worked to revamp the organization's digital presence by refining its media strategy across web and social platforms, addressing brand confusion with a similarly named national entity, and developing marketing templates and press outreach materials. Drawing on insights from a social media audit, they used AI tools to generate optimized content ideas and engagement strategies. They also identified relevant influencers and created customizable outreach prompts to expand the organization's reach.

### **Organization 3 | *Local History Preservation Project***

A team partnered with a long-running history preservation project to help transform its archive of oral histories, photographs, and historical documents into immersive, public-facing digital installations. Focused on three thematically distinct sites, the team aimed to lay the groundwork for a virtual walking tour that would make the community's rich history more accessible. AI tools

were used for tasks such as transcribing interviews, extracting themes, organizing archival material, and suggesting narrative structures. The final deliverables included foundational exhibit materials for each site and a set of practical insights into the benefits and constraints of using AI in historical storytelling.

#### **Organization 4 | *Cultural History and Local Tour Group***

A team collaborated with a local cultural tour organization that focuses on amplifying underrepresented histories in a Boston area neighborhood, particularly those of African American and Native American communities. The team was tasked with reviewing and improving the organization's educational curriculum by identifying historical gaps, structural inconsistencies, and citation issues. Using AI tools such as ChatGPT, Deepseek, and Claude, they analyzed the curriculum to address citation errors and benchmark the existing curriculum against Massachusetts educational standards. Deliverables included three detailed analysis reports, revised curriculum suggestions, and a set of hands-on AI prompts designed to help the partner experiment with integrating AI into their ongoing educational work.

#### **Organization 5 | *Nutrition-focused Advocacy Organization***

A team partnered with a nonprofit (more national in scope) focused on shifting public discourse around dairy consumption to support its advocacy for emerging policy/legislation that promotes milk-related product choice in schools. The project centered on using AI to strengthen public messaging, engage lawmakers and Congressional staff, and expand grassroots mobilization. The team developed AI-generated campaign content, including persuasive visuals and messaging for digital platforms, and used AI tools to analyze sentiment trends and legislative positioning. They also designed interactive chatbot prompts to test engagement strategies.

### **Lessons Learned**

The academic consulting teams proceeded according to action research principles, which involve iterative cycle phases of plan-act-observe-reflect. Often, they found that AI could provide tremendous value in terms of speed, precision, scope, and efficiency. Just as often, however, they realized that not all problems are "AI-amenable" problems – many issues facing civil society organizations are human-centered, and the solutions require complex, painstaking steps with human oversight and input.

The following are major general lessons learned, which were consolidated by the teams and then organized by the authors of this report. We also detail the technical contributions and lessons that relate to using the AI models as tools:

#### ***Team Technical Contributions***

Academic consulting teams helped translate the needs of civil society organizations into useful AI solutions. By combining technical skills with research and collaboration, they helped make the tools more effective and relevant.

Their key technical work included:

**1) Prompt Engineering and Few-Shot Learning:**

Graduate student consultants created structured prompt templates and added examples to help AI models produce better results. This worked well for tasks like rewriting text in a specific tone, summarizing documents, or editing drafts. In some cases, they also created dedicated prompt templates that civil society partners could reuse in their daily workflows.

**2) Creation of Custom GPTs and Tools:**

Teams built custom versions of GPTs or lightweight tools using APIs to automate tasks such as document formatting or creating images for social media handles or chatbot solutions for their websites.

**3) Data Cleaning and Preparation:**

When working with PDFs, spreadsheets, videos, or scanned files students used Python libraries to generate transcriptions, clean and organize the data so that AI models could understand it better.

**4) Comparing Model Performance & Identifying Best models for their needs:**

Teams used various platforms to test and compare the performance of different large language models (OpenAI GPT-4, Claude, Gemini, DeepSeek, etc.) on specific parts of their workflows and identified the best models for their needs.

**5) Collaborative Shared Workspaces:**

Teams made effective use of shared Claude projects and other collaborative tools to stay organized, track progress, and work together using a shared knowledge base.

### ***General Takeaways***

- **Learning before AI:** Before gaining extensive domain knowledge of the civil society organization by the academic consulting teams, AI tools were of much more limited in how well they could transform documents, extract relevant meaning, and search for additional resources. As teams gained more knowledge, they became more powerful users of the tools.
- **Research improves tool use:** Investing time and doing research before creating prompts really pays off. It is necessary to perform highly targeted work for civil society organizations that operate in very specific areas.
- **Stakeholder overestimation:** Because of varying levels of AI literacy and understanding, civil society organization stakeholders sometimes overestimated the capability of AI for

media and communications, suggesting the need for more AI training so capacity and capabilities are better understood. That said, academic consultant teams saw civil society stakeholders grow in their trust in models as careful mutual work and explanation were carried out.

- **Combinations for learning:** Civil society organizations often seek to engage their communities and audiences in learning activities. In such teaching/learning contexts, combinations of traditional education techniques, peer learning, and AI-assisted bots and tutoring may work best.

### ***AI Model Usage Lessons***

- Providing templates, transcripts, and example documents very much helps models with precision. Few-shot prompting – first providing the models with roughly analogous outputs sought by one’s prompts – is often necessary.
- Carefully choose the particular model to use for different tasks. In particular, the models seem to handle and analyze different document formats in diverging ways, with quality of outputs varying significantly. OpenAI’s ChaptGPT, Anthropic’s Claude, Google’s Gemini, Perplexity, and the Chinese model DeepSeek were the main LLM models that academic consulting teams experimented with, and they often compared model performance on the same tasks and inputs.
- For web retrieval tasks, the various models also vary significantly in their ability to identify targeted information, and quality can depend on the question at hand and the particular set of prompts. Experimentation is key, as is comparative use of models.
- For social media-related tasks, providing descriptions of what the civil society organization is and does, as well as content-specific prompts, seemed to help with relevance when generating accompanying visuals and art for posting.
- AI models can be helpful for message and marketing targeting, including finding aligned stakeholder individuals and groups, relevant influencers, and media organizations.
- AI models proved useful both in generating potential next-step action plans for social media calendars and, after combing through particular sets of social media posts, generating ideas for optimizing social media strategy.

### ***Organizational Legibility to AI***

A critical challenge identified across these consulting engagements relates to what we term "organizational legibility to AI"—the extent to which civil society organizations' unique contexts,

values, and operational realities can be effectively interpreted and acted upon by generative AI systems. Unlike large corporations with standardized processes and well-documented workflows, civil society organizations operate with idiosyncratic requirements rooted in specific missions, local contexts, and relationship-based communication strategies that are not immediately legible to AI models. This legibility concept we spotlight here is linked to the more general problem in AI research known as “alignment” (Shen et al., 2024), although we would argue that the problem of human-AI communication as applied in the civil society space takes on particular nuances that make it a distinctive area of study.

The consulting teams discovered that AI tools demonstrated limited effectiveness until significant domain knowledge about each organization was acquired and explicitly provided through carefully crafted prompts, templates, and contextual examples. This finding suggests that the “narrow” and “boutique” nature of nonprofit work—characterized by highly specific stakeholder groups, mission-first orientations, and community trust requirements—creates an inherent translation problem between organizational reality and AI comprehension. Without substantial human mediation to render organizational context, values, and intended outcomes legible to AI systems, these tools risk producing outputs that are technically proficient but fundamentally misaligned with the nuanced communication needs and authentic voice that define effective civil society engagement.

The experiences across these five case studies reveal a nuanced landscape where generative AI tools offer significant potential for civil society organizations, but only when deployed thoughtfully within existing organizational contexts and constraints. The most successful applications occurred when AI complemented rather than replaced human expertise and relationships—serving as an amplifier for domain knowledge rather than a substitute for it. Notably, the iterative nature of the consulting process itself proved as valuable as the final deliverables, suggesting that the “how” of AI integration may be as important as the “what” for mission-driven organizations operating with limited resources and high stakeholder trust requirements.

## **Recommendations**

The consulting engagements highlight several critical factors for successful AI integration in civil society contexts. Unlike large corporations with dedicated technical teams, these organizations require approaches that acknowledge their unique operational realities: mission-first orientations, relationship-based communication strategies, and resource limitations that demand high-impact, low-maintenance solutions. Based on the existing literature in this area and the experiences of the academic consulting teams, we also suggest several more general takeaways.

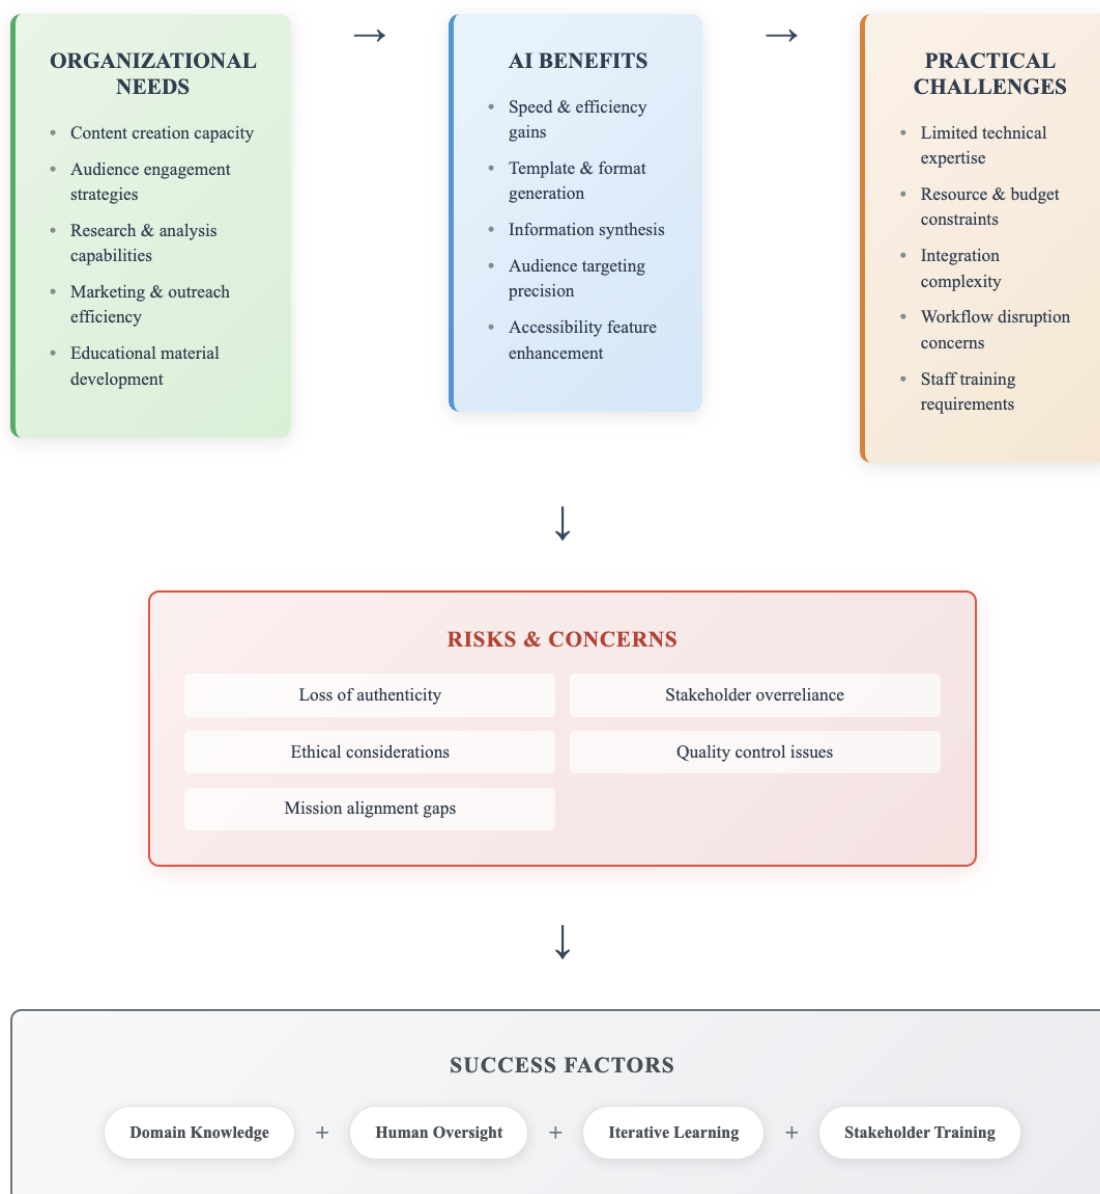
The following recommendations emerge from patterns observed across all five case studies and reflect both the potential and the constraints identified through the action research process.

- 1) **Invest in AI literacy before tool deployment:** Organizations should prioritize foundational training for staff and stakeholders to develop realistic expectations about AI capabilities and limitations, preventing both over-reliance and under-utilization of these tools.
- 2) **Start with documentation and templating tasks:** Begin AI integration with lower-risk activities like content organization, template creation, and research synthesis before moving to public-facing applications, allowing teams to build confidence and expertise gradually.
- 3) **Maintain human oversight for mission-critical communications:** Ensure all AI-generated content undergoes review by staff members with deep organizational and domain knowledge, particularly for materials that directly represent the organization's values and stakeholder relationships.
- 4) **Develop hybrid workflows that combine AI efficiency with authentic engagement:** Design processes that use AI for initial drafts, research, and organization while preserving space for human creativity, relationship-building, and contextual nuance that define effective civil society communication.
- 5) **Create sustainable, low-maintenance integration approaches:** Focus on AI applications that can be maintained with existing staff capacity and technical infrastructure, avoiding solutions that require ongoing specialized support or create new dependencies.

These recommendations reflect a pragmatic approach to AI adoption that recognizes both the transformative potential and the inherent limitations of current generative AI tools in mission-driven organizational contexts. The most effective integrations occurred when teams treated AI as a collaborator rather than a replacement, using these tools to enhance rather than automate the relational and contextual work that defines effective civil society communication. This suggests that successful AI adoption in these settings requires not just technical implementation, but also careful attention to organizational culture, stakeholder relationships, and the preservation of authentic voice and mission alignment.

A more general framework for understanding this problem space might usefully be rendered across four dimensions – needs, benefits, challenges, and risks. The following (Figure 1) is a preliminary visual scope of the problem space for AI consulting within the context of civil society organizations.

**Figure 1: Model of Civil Society AI Integration: Needs, Benefits, Risks, and Practical Challenges**



## Future Questions and Directions

Several important questions emerge from this preliminary exploration that warrant deeper investigation through more systematic research approaches. How do different types of civil society organizations—varying by size, mission focus, and community context—experience distinct benefits and challenges when integrating AI tools? The current case studies, while diverse, represent a limited sample from one metropolitan area, and future research should examine how

factors such as organizational age, funding structure, volunteer versus staff-driven operations, and community demographics influence AI adoption patterns. Additionally, longitudinal studies tracking these same organizations over extended periods would provide valuable insights into the sustainability of AI integration efforts and whether initial enthusiasm translates into lasting workflow changes or merely represents temporary experimentation.

The ethical dimensions of AI use in civil society contexts also require more nuanced examination, particularly regarding questions of representation, community voice, and power dynamics. How do AI-generated materials affect the authenticity and community trust that many civil society organizations depend on for their effectiveness? Future research should explore stakeholder perceptions of AI-assisted communications, examining whether community members can distinguish between human and AI-generated content and how such distinctions affect their engagement with organizations. Furthermore, as AI tools become more sophisticated and ubiquitous, there is a critical need to understand how their adoption might reshape the broader civil society landscape—potentially creating advantages for organizations with greater technical capacity while further marginalizing those without such resources, ultimately affecting the diversity and equity of voices in public discourse.

## References

- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... Liang, P. (2021). On the opportunities and risks of foundation models. arXiv. <https://arxiv.org/abs/2108.07258>
- Burt, E., & Taylor, J. A. (2003). New technologies, embedded values, and strategic change: Evidence from the U.K. voluntary sector. *Nonprofit and Voluntary Sector Quarterly*, 32(1), 115–127. <https://doi.org/10.1177/0899764002250006>
- Coghlan, D. (2011). AR: Exploring perspectives on a philosophy of practical knowing. *Academy of Management Annals*, 5(1), 53-87.
- DiSalvo, C., Clement, A., & Pipek, V. (2013). Participatory design for, with, and by communities. In J. Simonsen & T. Robertson (Eds.), *Routledge International Handbook of Participatory Design* (pp. 182–209). Routledge. <https://doi.org/10.4324/9780203108543>
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- Gálvez-Rodríguez, M. D. M., Caba-Pérez, C., & López-Godoy, M. (2016). Drivers of Twitter as a strategic communication tool for non-profit organizations. *Internet Research*, 26(5), 1052–1071.
- Guzman, A. L., & Lewis, S. C. (2024). What generative AI means for the media industries, and why it matters to study the collective consequences for advertising, journalism, and public relations. *Emerging Media*, 2(3), 347-355.
- Hussain, A., & Rizwan, R. (2024). Strategic AI adoption in SMEs: A prescriptive framework. arXiv. <http://arxiv.org/pdf/2408.11825.pdf>
- Light, A., & Miskelly, C. (2015). Sharing economy vs sharing cultures? Designing for social, economic and environmental good. *Interaction Design and Architecture(s) Journal*, 24, 49–62. [http://www.mifav.uniroma2.it/inevent/events/idea2010/doc/24\\_3.pdf](http://www.mifav.uniroma2.it/inevent/events/idea2010/doc/24_3.pdf)
- Lovejoy, K., & Saxton, G. D. (2012). Information, community, and action: How nonprofit organizations use social media. *Journal of Computer-Mediated Communication*, 17(3), 337–353. <https://doi.org/10.1111/j.1083-6101.2012.01576.x>
- O'Reilly Media. (2022). AI adoption in the enterprise 2022. <https://www.oreilly.com/radar/ai-adoption-in-the-enterprise-2022/>

Shen, H., Kneare, T., Ghosh, R., Alkie, K., Krishna, K., Liu, Y., ... & Jurgens, D. (2024). Towards bidirectional human-AI alignment: A systematic review for clarifications, framework, and future directions. <https://arxiv.org/abs/2406.09264>

Stringer, E. T. & Aragón, A. O. (2021). *Action Research* (5th ed.). Sage Publications.

West, D. M. (2018, January 4). The last mile problem in AI. Brookings Institution. <https://www.brookings.edu/articles/the-last-mile-problem-in-ai/>

Zhang, C., et al. (2023). The AI Index Report 2023. Stanford Institute for Human-Centered Artificial Intelligence. <https://aiindex.stanford.edu/report/>