

Building Interdisciplinarity in Engineering Doctoral Education: Insights from DTAIS Summer Incubator

Dr. Erica Cusi Wortham, GW Engineering

Erica Cusi Wortham is a cultural anthropologist with an interdisciplinary practice at George Washington University that spans social science, art design and engineering. As Director of the GW Engineering's Innovation Center | M06, she makes space for studio-based, tactile learning, community-driven innovation that centers creative problem-solving, equity and sustainability, and brings an ethnographic perspective to various AI initiatives at the school.

Prof. Zoe Szajnfarber

Dr. Zoe Szajnfarber is an Assistant Professor of Engineering Management and Systems Engineering at the George Washington University. Her research seeks to understand the fundamental dynamics of innovation in technology-intensive governmental organization, as a basis for decision-making. She received her bachelor's degree in Engineering Science from the University of Toronto. Szajnfarber conducted her graduate work at the Massachusetts Institute of Technology, earning dual masters' degrees in Aeronautics & Astronautics and Technology Policy and a doctorate in Engineering Systems. Her dissertation focused on technology infusion at NASA and involved substantial field work at the Goddard Space Flight Center. Outside of academia, Dr. Szajnfarber has worked as a systems engineer at MDA Space Missions (Canadarm Program) and Dynacon Inc. (Microsatellites); and as a researcher at the European Space Agency (Advanced Concepts Team).

Dr. Robert Pless, The George Washington University
Ryan Watkins, The George Washington University

Building Interdisciplinarity in Engineering Doctoral Education: Insights from DTAIS Summer Incubator

Abstract

In an era where the ethical, social, and technical complexities of artificial intelligence (AI) demand innovative approaches, the significance of diverse scholarship and interdisciplinary collaboration has never been more critical. Such collaboration is essential for the development of AI systems that are not only technologically advanced but also ethically sound and socially responsible. In 2021 GW Engineering was awarded funding to launch an interdisciplinary program on trustworthy artificial intelligence (AI). The *Designing Trustworthy AI in Systems* (or DTAIS) program brings together PhD students from the fields of Systems Engineering and Computer Science to co-design research, and tackle the conceptual and methodological bridge building that cross-disciplinary research and scholarship demands. This paper focuses on how this work has been accomplished thus far, in the context of the cornerstone Summer Incubator, and shares some of the lessons learned.

Introduction

The complexity of the ethical, social, and technical challenges associated with development trustworthy AI underscores the paramount importance of fostering diverse scholarship and interdisciplinary collaboration. It is through the melding of varied perspectives, expertise, and methodologies that we can create and manage AI systems that meet the diverse requirements for a system to be considered *trustworthy* [1]. These systems must embody ethical integrity and social responsibility, ensuring they adhere to the highest standards of fairness and accountability. Interdisciplinary collaboration, therefore, becomes not just beneficial but essential in the training of graduate students, enabling a synergy that drives the socio-technical development and application of AI technologies.

In 2021 GW Engineering was awarded funding to launch an interdisciplinary program on trustworthy AI [2]. *Designing Trustworthy AI in Systems* (or DTAIS) brings together PhD students from systems engineering and computer science to co-design research and tackle the conceptual and methodological bridge building that cross disciplinary work demands. The intent of the National Science Foundation Research Traineeship (NRT) program is to train STEM graduate students in “high priority interdisciplinary or convergent research areas, through a comprehensive traineeship model that is innovative, evidence-based, and aligned with changing workforce and research needs [3].” In what follows we share how the DTAIS program is designed to span interdisciplinarity and convergence with particular emphasis on a cornerstone Summer Incubator course. The 10-week Summer Incubator brings systems engineers and computer science PhD students in a modified design sprint to make sense of AI in real world settings and build short-run research prototypes together under the guidance of an interdisciplinary instructional team. The Summer Incubator helped establish a fertile middle ground where a mixed method ethos, design sprint rhythm and intentional sense of community enlivened (and complicated) the normative student-advisor modality most PhD students experience and fostered interdisciplinary collaboration. Some key takeaways from two iterations of the Summer Incubator course have to do with the benefits of 1) studio-based learning environments and the design framework and 2) building community intentionally. In line with the goal of NRT funding to create a sustainable, exportable, doctoral education program these experiences represent a potential roadmap for how to scaffold interdisciplinarity and convergence in engineering doctoral education.

Convergence research, one of the ten “Big Ideas” that drive NSF’s long-term agenda, is about merging ideas, approaches, and technologies to solve the world’s pressing problems. How to achieve convergence in research varies in practice; the DTAIS program maintains students in their discipline (their degree maps, undisturbed) and adds a series of courses that builds interdisciplinarity and seeds convergence through research practice around its theme, trustworthy AI. On a practical level, participants in the program develop depth in their core disciplines,

Computer Science or Systems Engineering, and become conversant in the non-core discipline. The Summer Incubator course (previously called “bootcamp”), which sits among an extended three-semester seminar series and a range of CS/SE courses (see Figure 1, DTAIS Timeline), is the program’s most significant pedagogical innovation.




Figure 1 DTAIS Timeline

Background

Algorithm designers today are not trained to think through the systems implications of their choices, and even more complicated, the societal implications [4], [5], [6]. At the same time, because of the pace with which AI innovations are affecting the world, trainees must excel not only in algorithmic design, mathematical rigor and programming patterns and abstractions but they must also be equipped to engage with the societal implications of their innovations [7], [8], [9], [10], [11], [24]. Similarly, while current systems programs are rapidly adopting advanced analytical techniques to inform systems analysis, the design aspect of the curriculum largely assumes persistence of existing system and process architectures [12], [13].

For algorithm designers, going forward this requires understanding, and being sensitive to, the context in which their creations will operate and evolve in unplanned ways through interaction with users in socio-technical ecosystems [14], [15], [16]; and for system designers, this means knowing enough about how AI tools are evolving to reimagine how tasks and processes can and should transform work in fundamental ways [17], [18], [19], [20]. In response the NRT students receive unique training preparing them for focused research projects that explore implications of bias in AI systems in general, and deeper dives into the tension between the opportunities and risks of integrating AI within organizations related to emergency response, smart and connected cities, and the cyber-social interaction landscape.

Incubator Design

The DTAIS program is designed with the goal of providing graduate students with intentional opportunities to develop their interdisciplinary scholarship in real-world contexts. The program does this through graduate seminars, collaborative courses, peer/faculty/industry mentoring, convenings/symposia/events, and the Summer Incubator. Within this framework, the primary goals of the incubator are to develop scholarly identity, build community, connect across disciplines, practice core research skills, learn ethics in context, and develop professional communication skills.

The Summer Incubator course combines a studio-based learning environment with a design framework whose built-in cycles of reflection and iteration – with an emphasis on prototyping – foster cross-disciplinary connections. We drew inspiration for the structure of the incubator from the design sprint [6], a method created to facilitate problem-solving or product launching within a compressed time frame, drawing out what is usually a five-day exercise over ten weeks. The rationale for this radical extension is a radical compression of what is often a year or more long transformation many PhD students undergo on their journey to a clear thesis proposal. Turning this rather solitary process into an intense and collective summer experience built around frequent interactions and iterations with feedback from peers, faculty and practitioners helped us model the bridge that links theory to practice and build a community of practice that includes diverse perspectives and is vertically integrated to include more senior students and faculty. We adapted and extended the four main sprint phases –map, sketch, prototype, and test – to allow time for reflection and for formulating research topics throughout the summer (See Figure 2, Summer Incubator Timeline, 2023).




Figure 2 Summer Incubator Timeline, 2023

During the mapping block students learned about pressing real-world needs through structured interactions with external partners and faculty. Following a brief overview about human-centered design and learning in real-world settings, what we call “AI in the wild,” the Summer Incubator course opens with a series of field trips and guest speakers. This focus on exposure engendered a discovery or curiosity mindset among the fellows; the cohort shared new experiences together and reflected aloud in the debrief on what they learned as well as on their disciplinary (and other) and biases. The cultural anthropologist on the instructional team brings an ethnographic approach into the mapping block. Getting students out of the classroom to observe or engage with social actors who use AI tools in their everyday life, in their jobs and companies, lends a grounded perspective to data provenance and the human and material infrastructures of AI that are often underappreciated. Data comes from people and places that are situated in complex contexts, and is always, by one means or another, collected via technologies or mechanisms that lay their own interpretive or computational (e.g. biased) fingerprint. Thus, to understand what constitutes trust and trustworthiness – an imperative that perhaps not unsurprisingly is still being widely discussed and debated in the DTAIS community – trust and trustworthiness must be *located* or identified in contexts where they are being achieved or sought. Trust and trustworthiness live under a variety of names, including at times in the negative, through its glaring absence with notion of harm and blame. Confidence, accuracy, usability, bias free, acceptable failure, explainability, transparency – these notions have all come up to give trust a situated or more nuanced definition in the “trust tracker” Google form students use to track how they understand the term.

The DTAIS program identifies three broad problem domains of interest that reflect research strengths of the participating departments: emergency management, cyber-social landscape, and smart, connected mobility. Summer Incubator fieldtrips were selected for their relevance to these domains. The first Summer Incubator took students to visit a rescue training camp and a private engineering research facility with active aviation and autonomous vehicle labs. In addition, executives and programmers from a telecom company came to speak with students on campus. For various reasons, the second Summer Incubator was structured somewhat differently with only one off campus field trip to a local grocery store outfitted with “Just Walk Out” technology and visits to GW robotics and drone labs. The first cohort of fellows were tapped to suggest AI in the wild experiences and through one student’s lead, we invited a senior data scientist from USAID to campus for an extended talk and discussion. A slightly deeper dive into two of the field trip experiences will convey the simple yet powerful reason for including opportunities to encounter “AI in the wild” during the Summer Incubator.

To begin, imagining the data landscape of a complex humanitarian disaster becomes a different conversation when you are staring down a dark, 30” diameter concrete tube piled under a mountain of building rubble. The Fairfax County Fire and Rescue facility was quiet the day we visited; walking the grounds with experienced rescue workers on a crystal-clear summer day belied the anxiety that percolated to the surface during the debrief session back on campus as fellows grappled with the responsibility of saving lives. The debrief sessions were conducted in an open-ended manner with a few tools that are perhaps more familiar to the designer than the engineer or computer scientist. Students were asked to sketch from memory an aspect of the visit that struck them. Most students drew the pile of rubble and the opening or tunnel presented by the cement tube. Sketches usually reveal perspective or positionality, whether intentional or not, that articulates something about how the fellow saw themselves. In this case, a flat, floor plan-like view, which was not actually experienced but used to make overall sense of what was seen and the spatial relationships among different assets on the training site contrasted several frontal views of the concrete tunnel. The tunnel, just large enough to fit a crouched average-sized human, glared as if from the belly of a beast that claims victims (See Figure 3, Student sketch from debrief session).



Figure 3 Student sketch from debrief session

Some fellows openly admitted that they were uncomfortable imagining their role as AI builders or model recommenders in life-or-death scenarios that would determine where to send a rescue team. How is trust achieved in this scenario? Quite differently, the visit to MITRE's autonomous vehicle and air traffic control labs gave the fellows an opportunity to see clearly demarcated and operationalized lines separating trust from trustworthiness. In a professional lab setting these demarcations formed the basis of a framework which supports the workflow and delineated roles and interactions where trust, or "calibrated trust" in this case, is looped with culture and people and trustworthiness, with competence and machines [21]. We entered murkier territory at the lab during a conversation with a retired aviator whose 35+ years of experience were being put to evaluate autonomous flight programs and protocols. The question of where and how humans remain *in* or *on* the loop in systems with AI tools continues to cast a long shadow among this research community; in fact, keeping track of the prepositions and their implications has become the subject of at least one research focus to emerge from the program. The 2023 field trip to Whole Foods' "Just Walk Out" store in Northwest D.C. put fellows directly in touch with store personnel and shoppers who manage or use the technology daily. In addition to the many humans who make the "Just Walk Out" experience feel seamless, constantly reshelfing and retagging items, the store itself drips with tech. Hundreds of cameras dot the ceiling and line each weighted shelf, tracking shoppers (and presumably workers) from entry to exit and every point in between. In this context AI tools feel suddenly tactile and the highly surveilled and regulated human labor on which they depend are thrown into relief, making at least some of the computational costs of

AI more visible [22]. “AI in the wild” experiences during the Summer Incubator serve to render the question of trust less abstract, forcing it to be located among real people who, depending on their roles, bring a range of agency and expertise to bear on the how the tools are designed and used. In addition, the experiences are designed to ask fellows as researchers to consider their own positionality in terms of “what does trust mean to you?” as we discuss below and in terms of where they sit within the systems boundaries they draw. The dense mesh of visible surveillance technology in the Whole Foods store made many of us confront our own selves as embodied data, training and engaging with AI tools that link the store to the online Amazon enterprise. The pile of concrete rubble in the emergency training facility aroused concern for life and death that framed trust with a gravity not fully considered prior to the visit.

Incubator Results

Debrief sessions used sketching for high-level impressions and takeaways followed by stakeholder mapping to draw system boundaries and understand who participates in each of the scenarios, including a consideration of what different roles, interests and influences are at play, and what problem statements (or research questions) are active in the space (associated with whose perspective). Identifying problem statements connected to actual people’s points of view is a key strategy of human centered design, facilitating the move from an abstract question to one that is grounded in context. Similarly, fellows were asked to map the technology they encountered in the site visits (and talks) over the stakeholders layer, adding depth and associations with the technology in play. Notes and observations from the field shared aloud became loops, dots and arrows that colored the whiteboards by the end of the debrief sessions, representations of collective work that reached across disciplines left out in the open as reference for subsequent inquiry. Perhaps of most use, the open conversation and mapping engendered conversations about trust with reference to real-world scenarios freshly experienced by the fellows. Trust is a construct; everyone agrees it does not exist inherently “out there,” nor can it be discovered in a positivist manner. As a construct it is achieved through interactions whether between actors, who often have different roles, training, experience, and access to information, or between actors and technologies. Plotting trust on the busy whiteboard in the form of a problem statement or research question makes evident that trust is a moving target and that it is up to the researcher to pin it down just long enough for it to be understood. Doing this work collectively, across disciplines using sketching and mapping techniques familiar to design studio settings, introduces a dose of humility (to balance disciplinary hubris) and lived experience while providing a fertile platform in which fellows become conversant in one another’s disciplines.

According to the external evaluator, the Summer Incubator course consistently and successfully deepened awareness of CS and SE related topics across disciplines, and among the positive effects of prototyping stand a solid awareness of how design choices impact implementation of AI tools in the field, and an established competency and confidence among the cohorts in leading

the development and adoption of trustworthy AI tools in works systems. Proficiency with non-technical skills such as communication, teamwork, and project management significantly improved over the course, with a resoundingly high score given to the value of learning the perspectives of diverse audiences. The intentional sense of community carefully built through shared experiences, reflection and necessary encouragement to try things enlivens the normative student-advisor modality with a strong interdisciplinary cohort.

Lessons Learned

The following blocks of the incubator – literature review, prototyping and testing – saw significant pivots from the first to the second summer iteration, in part because the program’s inaugural cohort was spread across very different stages of the PhD timeline. The external review process allowed the instructors and faculty advisors to home in on where to make improvements and, as a result, the two-week literature review block shifted from being fairly open, topically, to being relatively restricted to answering the broad question, “what is trust?” in their respective research domains. Attempts to do a community-based literature review the first summer was unsuccessful as fellows struggled with covering a wide range of readings on trustworthy AI that was a relatively new concentration for them. This block was also punctuated by several “soft skills” workshops led by a broader group of faculty, on how to prototype and later, policy, ethics and storytelling [23], [24], [25]. For the second Summer Incubator workshops were modulated into “broadening” talks with AI experts from other disciplines, such as philosophy and data science, which provided useful comparative points for discussion during literature review and helped keep the pace of the incubator a bit more manageable. The overt goal for the literature review was to support these scholars to arrive at confidently articulating a working definition of trust (or its corollary) relevant to their chose problem domain; the concrete expectation consisted of reading at least ten articles and arranging annotations from each article into a table that called out “fuzzy concept” (defining trust), conceptualization (how do they analyze trust), and operationalization (how do they measure trust). Drawing from the social sciences, tracing the move from fuzzy concept to conceptualization to operationalization reinforced the idea that arriving at one definition of trust is unrealistic and an undesirable goal; the point is to be able to talk about it with coherence, context specificity and complexity [26]. To aid in this move, fellows were asked to sketch answers to “what does trust mean to you?” using simple block diagramming that plots trust in a situated or contextual manner between the user and the algorithm, for example, or further out in the operation context, always with an awareness of vernacular terms that add a more nuanced understanding of trust as a construct (See Figure 4).

Society?




Figure 4 “What does trust mean to you?”

The prototyping block, which is a major emphasis of the incubator occupying four of the ten weeks, shifted dramatically from individual projects the first summer to a curated set of interdisciplinary team projects that were hammered out with faculty advisors months before the summer began. The reason for this change has to do with sorting out the fellow–faculty advisor relationship in a way that did not create competing interests for time and attention. Participation in the DTAIS program draws PhD students into a cohort while they maintain their disciplinary footing along with their prior and primary relationship with their advisors (and source of funding). Summer incubator projects have resulted in or built on several fellow–advisor co-authored papers and presentations [27], [28], [29] and one op-ed [30]. But to relieve the tension or lack of clarity from the first summer experience – where one fellow reported they did not have free time to pursue an independent project as their advisor expected and one advisor reported that the summer course should not be separate from the students’ research assistant commitments – the incubator team set about defining summer projects directly with the faculty advisors. In addition, fellows were teamed up in interdisciplinary pairs. This shift brought faculty advisors into the summer program in a meaningful way, alleviating conflicts in expectations with their students and making the occasional request for their presence throughout the summer seem less ornery. The interdisciplinary pairing of summer 2023 was in part modeled by a highly successful and organic inter-cohort collaboration that took off the previous summer and has resulted in a co-authored conference presentation on saliency maps [31].

The pre-decided summer projects also tackled a learned reluctance to try things out. PhD students are often given projects and they get started without ample time to think about the context and figuring out the right question to ask. The risk of failure – or of asking the wrong or less interesting question – is not well tolerated in academic settings, something the program tries to address by lowering the stakes and giving students time to experiment with the support of a community of peers and engaged faculty. Summer projects were not designed or destined to be core dissertation research projects; in fact, the idea was to develop skills in research formulation,

identifying problem statements practicing research skills, including identifying methodologies, communicating, giving and receiving feedback, presentation skills, navigating ethical dimensions – these were all given ample time via workshops and assignments in the first Summer Incubator and via peer and faculty advising in the second. The first Summer Incubator ran without pre-decided projects and some fellows spent a good portion of the four weeks figuring out what to try and, in some instances, had difficulty letting go of literature review. With projects already hammered out, the second summer’s prototyping block got off to a swift start with lively discussions on methods and data collection, setting up surveys and experiments. Fellows were asked to share and update prototyping plans with their advisors throughout the four-week prototyping block and work toward a presentation of their initial findings and experience.

The initial goal of the testing or validation block was to elicit substantive feedback from independent outsiders on the projects, with a focus on building communication skills through preparing and refining presentations before a month-long break in the summer schedule. The block started with how to give, hear and incorporate feedback and included a series of presentations that began with an audience of peers (including members of the earlier cohort who often attended incubator sessions) and extended to include DTAIS faculty. Presentations to external members of the community came once the fall semester began, in seminar. During the second summer incubator, the emphasis on presentations and feedback was modified to include more intentional, open reflection on each block of the incubator, with particular emphasis on what it means to be an interdisciplinary researcher in this space.

The next iteration of the Summer Incubator will run much like the second, though with a slightly more compressed timeline during the first block and a single pre-decided project that the fellows will work on together while tackling different aspects of the problem. Future iterations will seek to settle aspects of the Summer Incubator into summer research programs and seminars umbrellaed under GW’s newly launched *Trustworthy AI Initiative* (GW TAI) [32] while still maintaining a strong sense of community. Maintaining an interdisciplinary research community requires that the connections and community aspects fostered during DTAIS continue to be nurtured as the initiative grows and new courses are charted to navigate the incredibly fast-paced and dynamic arena of trustworthy AI.

Conclusions and Recommendations

The DTAIS Summer Incubator represents a pioneering endeavor in engineering doctoral education, emphasizing the importance of interdisciplinarity and convergence in the development of trustworthy AI systems. Through its innovative approach –melding systems engineering with computer science – the program is successfully fostering a unique educational environment. This environment not only cultivates technical skills but also nurtures a deep understanding of the ethical, social, and contextual dimensions of AI technology. The incubator's studio-based

learning, combined with its focus on real-world applications, prototyping, and community building, has proven to be a robust model for encouraging cross-disciplinary collaboration and innovation.

The iterative nature of the program, highlighted by adaptations and refinements across its iterations, underscores the value of flexibility and responsiveness in educational program design. By integrating feedback from participants, faculty, and external evaluators, the Summer Incubator has evolved to better meet the needs of its students and to more effectively achieve its objectives of developing scholarly identity, building community, and preparing students for the challenges of designing AI within complex socio-technical systems.

For institutions considering the implementation of similar interdisciplinary programs, we offer the following recommendations based on our experiences with the Summer Incubator:

Embrace Studio-Based Learning: A studio-based learning environment, characterized by its emphasis on collaboration, prototyping, and iterative design, is crucial for fostering creativity and innovation. Encourage students to engage in hands-on projects that address real-world problems, facilitating a deeper understanding of the complexities involved in trustworthy AI development.

Facilitate Real-World Exposure: Incorporating field trips and guest lectures from industry experts provides students with invaluable insights into the application of AI technologies in various sectors. This exposure is essential for broadening students' perspectives and enhancing their ability to identify and tackle relevant challenges within their research.

Prioritize Community Building: The development of a supportive and inclusive community is key to the success of interdisciplinary programs. Activities that encourage reflection, shared experiences, and open dialogue can help to forge strong bonds among participants, fostering a collaborative environment conducive to interdisciplinary learning.

Iterate and Adapt: Be prepared to iteratively refine the program based on feedback from all stakeholders, including students, faculty, and industry partners. Flexibility in program design allows for the incorporation of new insights and the addressing of challenges as they arise, ensuring the program remains relevant and effective.

Integrate Ethics and Societal Impact: Ensure that discussions on ethics and the societal impact of AI are woven throughout the curriculum. Preparing students to consider the broader implications of their work is critical for the development of trustworthy AI systems that are not only technologically advanced but also ethically sound and socially responsible.

By applying these recommendations, institutions can enhance their ability to prepare the next generation of engineers and computer scientists for the challenges of designing AI systems that are both innovative and trustworthy.

- [1] "Trustworthy and Responsible AI," *NIST*, Jul. 2022, Accessed: Apr. 07, 2024. [Online]. Available: <https://www.nist.gov/trustworthy-and-responsible-ai>
- [2] National Science Foundation, *Award Number 2125677*.
- [3] National Science Foundation, [Online]. Available: <https://new.nsf.gov/funding/opportunities/national-science-foundation-research-traineeship>
- [4] K. Martin, "Ethical Implications and Accountability of Algorithms," *J Bus Ethics*, vol. 160, no. 4, pp. 835–850, Dec. 2019, doi: 10.1007/s10551-018-3921-3.
- [5] S. Newell and M. Marabelli, "Strategic Opportunities (and Challenges) of Algorithmic Decision-Making: A Call for Action on the Long-Term Societal Effects of 'Datification,'" *SSRN Journal*, 2015, doi: 10.2139/ssrn.2644093.
- [6] S. Faraj, S. Pachidi, and K. Sayegh, "Working and organizing in the age of the learning algorithm," *Information and Organization*, vol. 28, no. 1, pp. 62–70, Mar. 2018, doi: 10.1016/j.infoandorg.2018.02.005.
- [7] P. Singh, "Beyond the basics: Exploring the impact of social media marketing enablers on business success," *Heliyon*, vol. 10, no. 5, p. e26435, Mar. 2024, doi: 10.1016/j.heliyon.2024.e26435.
- [8] I. M. Cockburn, R. Henderson, and S. Stern, *The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis*. in NBER book chapter series, no. no. c14006. Cambridge, Mass: National Bureau of Economic Research, 2018.
- [9] G. Eden, M. Jirotko, and B. Stahl, "Responsible research and innovation: Critical reflection into the potential social consequences of ICT," presented at the IEEE 7th International Conference on Research Challenges in Information Science (RCIS), 2013, pp. 1–12.
- [10] J. L. King, V. Gurbaxani, K. L. Kraemer, F. W. McFarlan, K. S. Raman, and C. S. Yap, "Institutional Factors in Information Technology Innovation," *Information Systems Research*, vol. 5, no. 2, pp. 139–169, Jun. 1994, doi: 10.1287/isre.5.2.139.
- [11] J. Kleinberg, "Inherent Trade-Offs in Algorithmic Fairness," in *Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems*, Irvine CA USA: ACM, Jun. 2018, pp. 40–40. doi: 10.1145/3219617.3219634.
- [12] N. Gislason, "Architectural design and the learning environment: A framework for school design research," *Learning Environ Res*, vol. 13, no. 2, pp. 127–145, Jul. 2010, doi: 10.1007/s10984-010-9071-x.
- [13] C. L. Dym, A. M. Agogino, O. Eris, D. D. Frey, and L. J. Leifer, "Engineering Design Thinking, Teaching, and Learning," *Journal of Engineering Education*, vol. 94, no. 1, pp. 103–120, Jan. 2005, doi: 10.1002/j.2168-9830.2005.tb00832.x.
- [14] P. M. Bednar and C. Welch, "Socio-Technical Perspectives on Smart Working: Creating Meaningful and Sustainable Systems," *Inf Syst Front*, vol. 22, no. 2, pp. 281–298, Apr. 2020, doi: 10.1007/s10796-019-09921-1.
- [15] R. Holton and R. Boyd, "'Where are the people? What are they doing? Why are they doing it?' (Mindell) Situating artificial intelligence within a socio-technical framework," *Journal of Sociology*, vol. 57, no. 2, pp. 179–195, Jun. 2021, doi: 10.1177/1440783319873046.
- [16] F. W. Geels, *Technological transitions and system innovations: a co-evolutionary and socio-technical analysis*. Cheltenham, UK ; Northampton, MA: Edward Elgar, 2005.
- [17] P. R. Daugherty and H. J. Wilson, *Human + machine: reimagining work in the age of AI*. Boston, Massachusetts: Harvard Business Review Press, 2018.
- [18] D. M. West, *The future of work: robots, AI, and automation*. Washington, D.C: Brookings Institution Press, 2018.

- [19] M. Meacham, *AI in talent development: capitalize on the AI revolution to transform the way you work, learn, and live*. Alexandria, VA: ATD Press, 2021.
- [20] P. R. Daugherty and Jim Euchner, "Human Machine: Collaboration in the Age of AI: Paul Daugherty talks with Jim Euchner about a new paradigm for collaborative work between people and intelligent systems," *Research-Technology Management*, vol. 63, no. 2, pp. 12–17, 2020.
- [21] Lacher, Andrew, "A Framework for Discussing Trust in Increasingly Autonomous Systems," The MITRE Corporation, Jun. 2017.
- [22] Kate Crawford, *Atlas of AI: Power, Politics and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.
- [23] K. Lindvig, "The implied PhD student of interdisciplinary research projects within monodisciplinary structures," *Higher Education Research & Development*, vol. 37, no. 6, pp. 1171–1185, Sep. 2018, doi: 10.1080/07294360.2018.1474343.
- [24] M. E. Exter, C. M. Gray, and T. M. Fernandez, "Conceptions of design by transdisciplinary educators: disciplinary background and pedagogical engagement," *Int J Technol Des Educ*, vol. 30, no. 4, pp. 777–798, Sep. 2020, doi: 10.1007/s10798-019-09520-w.
- [25] A. Buhmann and C. Fieseler, "Towards a deliberative framework for responsible innovation in artificial intelligence," *Technology in Society*, vol. 64, p. 101475, Feb. 2021, doi: 10.1016/j.techsoc.2020.101475.
- [26] A. Bhattacharjee, *Social science research: principles, methods, and practices*, Second edition. Tampa, Florida: Anol Bhattacharjee, 2012.
- [27] Zaid Kbah and Erica Gralla, "Understanding Enablers and Barriers for Deploying AI/ML in Humanitarian Organizations: the case of DRC's Foresight," presented at the Institute of Industrial and Systems Engineers (IISE), New Orleans, LA, May 2023.
- [28] Aditya Singh and Zoe Szajnfarder, "Comparative Analysis of Pathways to Changeability," in *Twentieth Annual Acquisition Research Symposium*, Naval Postgraduate School, Monterey CA, Apr. 2023, pp. 148–156.
- [29] Aditya Singh and Zoe Szajnfarder, "Understanding Post-Production Change and Its Implications for System Design: A Case Study in Close Air Support During Desert Storm," *Naval Engineers Journal*, vol. 134, no. 3, pp. 87–95, 2022.
- [30] Haneen Al-Rashid, "AI in global development is more than just a set of tools," *Devex.com*, Jan. 04, 2024. [Online]. Available: <https://www.devex.com/news/opinion-ai-in-global-development-is-more-than-just-a-set-of-tools-106853>
- [31] Chris Krueger, Justine-Louise Manning, Robert Pless, and Zoe Szajnfarder, "A Saliency Map is Worth a Thousand Words," presented at the Military Operations Research Symposium, West Point NY, 2023.
- [32] "Leading the Way to AI You Can Trust," *GW Today*, Apr. 02, 2024. [Online]. Available: <https://gwtoday.gwu.edu/leading-way-ai-you-can-trust>