AI:XPERTISE Lab: Embedding AI in Expert Work

Andreas Møgelmose *Dept. of Architecture, Design and Media Technology, Aalborg University* anmo@create.aau.dk
Kasper Trolle Elmholt *Dept. of Politics and Society, Aalborg University* elm@dps.aau.dk

Abstract

Advanced AI technologies increasingly encroach on domains traditionally occupied by expert professions such as medicine and architecture. Existing work often isolates technical development from the organisational conditions that shape how AI is adopted. AI:XPERTISE is a new lab and research project at Aalborg University, which brings these perspectives together by co-designing new AI systems while simultaneously studying their integration into expert workflows. This position paper outlines the lab's research agenda, methodological approach, and expected contributions.

Keywords: AI, expert work, sociotechnical systems

1 Introduction

Artificial intelligence is frequently promoted as a solution capable of transforming professional work (Baer et al. [2025]). Reports suggest that up to seventy percent of tasks may be automated by 2030 (Durth et al. [2024]). Field experiments additionally show that while generative tools can boost knowledge worker productivity, quality outcomes vary (Dell'Acqua et al. [2023]). Yet real-world deployments remain fraught with difficulty. Empirical studies show that many AI projects stall or are abandoned before delivering value (Lebovitz et al. [2021]). Scholars increasingly argue that focusing solely on algorithmic performance obscures the organisational context in which systems are developed and used (Bailey and Barley [2020]). Professionals also grapple with opaque model recommendations when using AI for critical judgements, underscoring the need to understand augmentation rather than substitution (Lebovitz et al. [2022]).

This paper introduces *AI:XPERTISE*, a new laboratory at Aalborg University dedicated to investigating the interplay between design of AI systems and professional expertise. Rather than treating technology and organisation as separate spheres, the lab views them as co-constitutive. We argue that successful AI adoption requires insights from computer science and organisational studies, both augmented by domain-specific knowledge. The lab therefore brings together researchers from engineering and the social sciences, supported by practitioners and policymakers, to explore how AI can be responsibly embedded in expert work.

2 Lab Vision and Background

For a while it was thought that AI would first automate blue collar jobs, but currently it seems that advanced AI technologies are much more suited for use in domains traditionally occupied by expert professions such as doctors, architects, and lawyers. While politicians and think tanks highlight the potential for job savings (Folketinget [2022], Durth et al. [2024]), considerable research shows persistent challenges when AI is introduced into expert work (Lebovitz et al. [2022], Baer et al. [2025]). AI:XPERTISE seeks to address this tension by developing AI systems in tandem with organisational analysis, ensuring that technical innovation aligns with professional values and institutional constraints.

The lab's guiding research question is: *How can AI technologies be successfully developed and embedded into professional expert work?* To answer this question we analyse contrasting sectors that exhibit different risk profiles, regulatory regimes, and labor practices. In architecture, we explore generative models that augment the visualisation capabilities of architects. In medicine, we investigate both large language models and computer vision models in contexts of anaesthesiology and surgery. These cases allow us to compare private-sector, low-risk environments with high-stakes public-sector settings. Additionally, architecture firms are traditionally very agile and technology-ready, while hospitals move more slowly.

AI:XPERTISE positions itself as a hub for interdisciplinary research on AI and expert work. The lab combines state-of-the-art techniques in generative and traditional AI modelling with insights from organisational theory and sociology of professions. By building and testing systems in real workplaces, we aim to uncover how technical design choices interact with organisational structures and cultural norms. The lab will also serve as a space for debate and reflection about the broader societal implications of AI, hosting seminars with academics, practitioners, and policymakers.

3 Related Work

Research on AI adoption in professional domains often focuses on either technical challenges or organisational factors, leaving a gap at their intersection. Studies of machine learning development highlight the difficulty of eliciting expertise and maintaining data pipelines (de Souza Nascimento et al. [2019]). Work in information systems shows that black-boxed models can disrupt established divisions of labour and decision rights (Lebovitz et al. [2021], Hillebrand et al. [2025]). Meanwhile, technical communities explore methods such as physics-informed learning and human-in-the-loop design to improve model reliability and transparency (Karniadakis et al. [2021], Bailey and Barley [2020]). Our lab builds on this literature by empirically linking design decisions to organisational outcomes across multiple domains, offering a comparative perspective that is currently missing.

4 Research Methodology

Our methodological approach integrates technical experimentation with qualitative inquiry. Each domain—architecture and medicine—will host a pair of PhD projects. One student will pursue a technical orientation, engaging directly with model architecture, data curation, and evaluation metrics. The second student will adopt a social science perspective, mapping organisational contexts, work practices, and governance structures. Together they will conduct iterative cycles of design, deployment, and reflection. The students will be supervised by an associate professor from their field (AI or organizational studies), and co-supervised by an associate professor from the opposite.

The architecture case focuses on generative AI for visualising building designs. We will experiment with models that combine data-driven learning with rule-based constraints such as physical laws or design guidelines, drawing on advances in physics-informed machine learning (Karniadakis et al. [2021]). The research investigates how to elicit tacit knowledge from architects, how to incorporate aesthetic criteria into model training, and how generative tools influence division of labour in design studios.

The medical case examines large language models and computer vision for supporting pre-anaesthetic screenings, as well as computer vision for situational understanding in surgery. We study both model development, techniques for evaluating model decisions, as well as the organisational procedures required to integrate AI into clinical workflows. Ethical considerations, patient safety, and data governance are central topics, especially as clinicians must interpret and occasionally contest AI outputs when stakes are high (Lebovitz et al. [2022]).

Across both domains we employ an experimental comparative case methodology (Yin [2009]). This approach enables systematic comparison of how similar technologies fare in divergent organisational settings. Data will be collected through ethnographic fieldwork, interviews, and system logs. Access to expert practitioners and datasets is secured through partnerships with architecture firms and Rigshospitalet. The research design is informed by literature on machine learning development challenges (de Souza Nascimento et al. [2019]) and organisational change in the face of AI (Hillebrand et al. [2025]).

5 Expected Outcomes and Impact

AI:XPERTISE aims to deliver both theoretical and practical contributions. First, we will develop a sociotechnical framework that theorises the conditions under which AI can enhance expert work. This framework will articulate how factors such as task structure, organisational governance, and model interpretability interact to support or hinder adoption. Second, we will produce a scalable design and implementation framework for practitioners. By distilling lessons from architecture and medicine, we will outline guidelines for eliciting expertise, configuring human-AI collaboration, and evaluating outcomes when ground truth is uncertain (Lebovitz et al. [2021]).

The project is structured around key milestones. By the end of 2025 we will map the organisational contexts and establish baseline technologies for both domains. By September 2026 we will start to evaluate the prototypes in real settings and document their effects on work practices. This will mark the start of iterative improvements which will run throughout the 3 year study. By July 2028 we expect to finalise the sociotechnical framework and disseminate design guidelines. The team aims to publish 12–15 peer-reviewed articles across computer science, information systems, and organisational journals and conferences.

Beyond academia, AI:XPERTISE seeks to influence public discourse on AI and labour. Political ambitions in Denmark call for saving thousands of public-sector jobs through automation (Folketinget [2022]). Our research will critically assess such expectations by documenting the practical challenges and resource implications of deploying AI in expert settings. The lab will also create resources for industry and public-sector partners, including workshops, open datasets where possible, and a community of practice around responsible AI.

6 Team and Collaboration

The lab is led by Principal Investigators Andreas Møgelmose and Kasper Trolle Elmholt, representing engineering and social science perspectives respectively. Both have experience with interdisciplinary collaborations and will coordinate the dual focus on technical rigour and organisational relevance. Professor Thomas B. Moeslund serves as mentor, drawing on extensive expertise in computer vision and AI in society via his involvement in Center for AI in Society (CAISA).

Two PhD students in each domain will be co-located and rotate between departments to foster knowledge exchange. Regular cross-departmental meetings ensure that insights from organisational analysis inform technical design decisions and vice versa.

To build a broader community, the lab will host seminars, inviting practitioners, policymakers, and academics to discuss emerging results. These events will function as forums for disseminating tools, sharing datasets, and exploring collaborative projects. By creating spaces where technical and organisational expertise intersect, AI:XPERTISE aims to nurture a new generation of researchers capable of navigating the complexities of AI in professional work.

7 Conclusion

AI:XPERTISE addresses the gap between technical innovation and organisational readiness. By jointly designing AI systems and examining their adoption in expert work, the lab will contribute theoretical insights, practical frameworks, and a collaborative community focused on responsible AI integration. We invite researchers and practitioners to engage with the project as we explore how AI can augment expert labour rather than replace it.

References

I. Baer, L. Waardenburg, and M. Huysman. What is augmented? a meta-narrative review of AI-based augmentation. *Journal of the Association for Information Systems*, 26(3):760–798, 2025.

D. E. Bailey and S. R. Barley. Beyond design and use: How scholars should study intelligent technologies. *Information and Organization*, 2020.

E. de Souza Nascimento et al. Understanding development process of machine learning systems: Challenges and solutions. In *IEEE International Symposium on Empirical Software Engineering and Measurement*, 2019.

F. Dell'Acqua et al. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Technical Report Working Paper 24-013, Harvard Business School Technology & Operations Management Unit, 2023.

S. Durth et al. The organization of the future: Enabled by gen AI, driven by people. Report, 2024.

Folketinget. Høring om kunstig intelligens og digitaliseringens betydning for fremtidens arbejdsmarked. Parliamentary hearing report, 2022.

L. Hillebrand, P. Andersen, and K. Petersen. Managing with artificial intelligence. *Academy of Management Annals*, pages 343–375, 2025.

G. E. Karniadakis et al. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

S. Lebovitz, H. Lifshitz-Assaf, and N. Levina. Is AI ground truth really true? the dangers of training and evaluating AI tools based on experts' know-what. *MIS Quarterly*, 45(3):1501–1526, 2021.

S. Lebovitz, H. Lifshitz-Assaf, and N. Levina. To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1):126–148, 2022.

R. K. Yin. *Case Study Research: Design and Methods*. Sage, 2009.