Proceedings of the Design Society, Volume 6: DESIGN 2026
https://doi.org/10.1017/pds.2026.10373

Boundary objects in co-designing AI systems: a case study of supporting cross-functional collaboration

Álvaro Aranda-Muñoz[1,✉], Yuji Yamamoto[2] and Kristian Sandström[2]

[1] Research Institutes of Sweden, Sweden, [2] Mälardalen University, Sweden

✉ alvaro.aranda.munoz@ri.se

ABSTRACT: In the early design stage of AI systems, designers face the challenge of addressing interdisciplinary needs with technical feasibility. This case study presents learnings from engaging a cross-functional team in the early design of an AI Vision System to support operators in the assembly process of electrical equipment. Our approach has been to embed Human-AI guidelines into boundary objects. The findings indicate that boundary objects helped participants reach common understandings, identify needs, and develop plans to mitigate future problems during the development of the AI system.

KEYWORDS: collaborative design, case study, artificial intelligence (AI), early design phase, boundary object

1. Introduction

The integration of AI systems into organisational workflows presents various challenges that often lead to implementation failure. Approximately 90% of projects fail at this stage (Yildirim, 2024, Chapter 1), particularly when systems must be integrated into work-oriented practices (Cai et al., 2019; Yang, Steinfeld, et al., 2019). This high failure rate stems partly from a disconnect: data scientists often create solutions that fail to address user needs, while designers are insufficiently involved in problem formulation to shape those needs effectively (Yildirim, 2024, Chapter 1). More broadly, implementation is often carried out without careful consideration of the employees who will work alongside AI systems (Makarius et al., 2020). Addressing this effectively requires, among other things, establishing a shared AI vision and enabling effective collaboration between different roles in the organisation (Haefner et al., 2023). To address these challenges, various resources (Jansen & Colombo, 2023; Yildirim, Oh, et al., 2023), processes (Colombo & Costa, 2024), and guidelines have been proposed by researchers (Amershi et al., 2019) and by technology companies such as Google (PAIR, 2021). Moreover, although recent literature has begun to examine how practitioners use these guidelines (Yildirim, Pushkarna, et al., 2023), research on their practical use and benefits remains scarce (Yildirim, 2024).

These deployment challenges are amplified by the complexity of AI as a design material. AI presents numerous challenges along the different stages of the design process (Yang, 2021). These challenges include difficulties understanding AI capabilities, assessing feasibility given available data, and the lack of boundary objects that facilitate cross-collaboration among designers, engineers, and other stakeholders (Yang, 2021, p. 18). Such cross-collaborations are difficult to achieve due to the absence of shared language, communication, and workflows (Nahar et al., 2022). Boundary objects offer a potential solution to these collaboration challenges. The concept of boundary object is a key aspect of collaboration among different communities of practice (Star & Griesemer, 1989).

Previous literature underscores how boundary objects can support cross-functional teams in assessing ML models, discussing end-user requirements, and explaining AI concepts to non-experts (Ayobi et al., 2021; Cai et al., 2019, 2021; Yang, Cranshaw, et al., 2019). For instance, Cai et al. (2019) applied the concept of boundary objects to create human-AI onboarding materials that introduced clinicians to a

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

decision support system. Their case study revealed that participants valued specific information about the system, including its capabilities and limitations, functionality, the system's perspective, optimisation goals, and considerations necessary before adoption.

However, despite progress in understanding boundary objects and their applications, important gaps persist. First, there is limited documentation of how practitioners adapt Human-AI (HAI) guidelines to specific contexts (Yildirim, 2024, p. 9). Second, while recent research has underscored the need for holistic AI development frameworks in manufacturing (Yamamoto et al., 2025), empirical work on the tangible resources required to implement them remains scarce. Specifically, resources that bridge the "separation of concerns" (Subramonyam et al., 2022) between diverse experts remain under-documented in manufacturing, despite a growing body of literature demonstrating the value of boundary objects in cross-functional teams in other domains (Ayobi et al., 2021; Cai et al., 2019, 2021; Yang, Cranshaw, et al., 2019). Addressing these gaps is essential to guide design practitioners in creating materials that onboard cross-functional stakeholders and facilitate early discussions on human-AI guidelines, leading to feasible, human-centred AI integrations. This paper addresses these gaps through a case study that examines and documents how boundary objects facilitated cross-functional collaboration during the early design and development of an AI quality-control system that uses real-time image recognition in a manufacturing environment in Sweden. The research team aimed to support the exploration of integrating this AI system into the assembly line by facilitating collaboration among assembly operators, production leaders, quality technicians, data scientists, and other stakeholders. To guide the case study, we present the following two research questions:

- RQ1: How can designers embed Human-AI guidelines into boundary objects to support cross-functional teams in the early design and development stages of an AI system?
- RQ2: How do boundary objects support different stakeholders in cross-functional teams during the early design and development of an AI system?

2. Background of the case study

This case study is part of a three-year research project aimed at democratizing AI adoption in manufacturing through a collaborative methodology that addresses human needs, involving five industrial manufacturing companies in Sweden. The research team (see Method) collaborated with a large electric power infrastructure manufacturer. This manufacturer had previously collaborated with the research team for one year on an earlier case involving an AI model for detecting temperature anomalies in furnace sensors (Yamamoto et al., 2024), during which we identified key challenges in early-phase AI system design and the need for more actionable methodologies in manufacturing contexts. While this work led to a systematic literature review and a broader development framework for human-AI systems in manufacturing (Yamamoto et al., 2025) containing 147 task recommendations across six stages, this case study focuses on how boundary objects help cross-functional collaboration during early design phases, underscoring the design perspective and the role designers can play in facilitating collaboration.

A two-hour factory tour revealed that electric equipment assembly is entirely manual and is described as artisanal, relying heavily on operator experience, knowledge, and skills. Operators are not considered experienced until they have worked six to seven years on the assembly line. Assembly varies with equipment size and customer specifications, ranging from a few days (across three shifts) to several weeks. Operators follow detailed, colour-coded instructions displayed on eye-level screens next to the assembly area. The process is iterative and circular, with interconnected layers in which some steps affect both past and future work. This complexity means that assembly errors can become hidden between layers, making human errors both common and costly if not detected early.

To minimise errors, operators conduct manual quality inspections with a mentor operator after each shift, reviewing all completed work. When errors are detected, responses vary in severity, ranging from later-stage corrections to undoing work until it matches the instructions. Undetected errors identified only during testing are more costly: assembly defects can cause equipment malfunctions and, in extreme cases, require discarding the equipment and restarting the entire work order. The tour concluded in the testing department, highlighting the high costs of late-detection errors.

Following the factory tour, the contact person and a data scientist presented a concept for the AI system, along with a machine-learning prototype, to demonstrate feasibility. The AI system aimed to address these challenges by helping operators detect assembly errors early, thereby reducing rework and alleviating the burden of costly mistakes through augmentation and assistance in manual assembly.

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

3. Method and data collection

The empirical material consisted of field notes taken during the factory tour, documentation handed by the case company informing about the AI system (e.g., information about the idea, the prototype, the estimated savings in costs), audio recordings of the three workshops, video recordings of presentations made by participants (in the three workshops), photographs and notes taken during the workshops, participants' feedback and reflections about the method and materials (boundary objects) at the end of the three workshops (audio-recorded), and the resulting material co-created by the participants. The research team conducted a qualitative analysis of this material through reflective sessions to examine the role of boundary objects in the workshops and their support for the cross-functional team.

The research team conducting this case study comprised four members: a PhD in human-computer interaction (HCI) with a computer science background, a senior lecturer in lean production, a senior production consultant, and a computer science professor. The workshop participants (identified in the first workshop) consisted of multiple organizational roles: a project leader in R&D (our contact person), the head of assembly, four assembly operators (three experienced mentors with more than six years of experience and one regular operator), one production technician, two quality technicians, a data scientist, and one R&D manager observing to learn the approach for future projects.

4. Materialising human-AI guidelines into boundary objects

The researchers had experience applying HAI guidelines from a previous case (Yamamoto et al., 2024), where they experienced some challenges when applying directly the PAIR guidelines (PAIR, 2021) due to difficulties in addressing operational workflows, the uncertainty of the model outputs, the interdependency between the AI model, the data and the user needs, difficulties to prioritize guidelines and a lack of a holistic view in the cross-functional team (Yamamoto et al., 2024). Some of these challenges have been observed in interviews with practitioners working with AI, who need to adapt guidelines to their domains and specific contexts (Yildirim, Pushkarna, et al., 2023).

To mitigate such challenges, for the case presented in this paper, the researchers drew upon best practices previously identified concerning the use of boundary objects in cross-functional teams with AI (Ayobi et al., 2021; Cai et al., 2019, 2021; Yang, Cranshaw, et al., 2019), considering HAI guidelines (PAIR, 2021) and previous research on guidelines (Amershi et al., 2019; Subramonyam et al., 2022) as a checklist of key aspects to address when engaging the cross-functional team. The researchers filtered the guidelines for use in the manufacturing context, taking into account contextual information obtained from the factory tour and meetings with various stakeholders. Then, the researchers shaped the resulting guidelines into boundary objects that could be used in workshop activities to help reach shared understandings of the potential AI system and to support early discussions on end-user needs, integration into the assembly line, the ML model, and potential data requirements.

In the following subsections, we describe the three-workshop journey, underscoring how each boundary object used reflected the guidelines and our perceived impact on the support these boundary objects provided to the workshops. Table 1, in the results section, summarises and illustrates these relationships.

4.1. Workshop 1: onboarding the cross-functional collaboration

The first workshop aimed to identify and establish the cross-functional team involved in the early design and development of the AI system. It began with a brief introduction to the AI system and the project's mission statement. Then, the researchers identified the components necessary for developing an AI application, drawing on the work of Subramonyam et al. (2022), namely the user's mental model, the user interface, the AI models, and the training data. The researchers underscored how vital it is to attend to contextual aspects to have successful integrations, reminding participants of an example previously explored in the research project from the work of Yang et al. (2019), and by introducing the topic of swimlanes to discuss workflows in the integration of AI systems (Sendak et al., 2020). Then, the researchers introduced a swimlane diagram template that participants completed, specifying which roles, both within and outside the assembly line, would be required to successfully address AI system alarms. This activity helped identify key roles and develop an initial understanding of the number of people who would be involved in and influenced by the AI system. The next activity of the workshop involved discussing the aspects necessary for a first minimum viable product that could be tested as a pilot on the

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

assembly line. To that end, the researchers developed a template to identify AI development-related needs, drawing on Subramonyam et al.'s (2022) work. Lastly, the researchers facilitated follow-up activities to prioritise and a fifteen-minute reflection.

4.2. Workshop 2: collaborative exploration of the AI system opportunity space

The second workshop lasted approximately two and a half hours. It consisted of onboarding activities to introduce the project, AI vision technologies, a presentation of the idea for the case with a demonstration of the currently developed prototype, and two activities to discuss initial aspects of the vision system and facilitate early discussions about potential integration aspects of such a support tool on the assembly floor. The workshop concluded with a fifteen-minute reflective activity.

In each activity, different boundary objects were employed to facilitate shared understandings and discussion. First, in the introduction to AI vision tools, we demonstrated Teachable Machine (Carney et al., 2020) to enable participants to explore, experience, test, and reflect on the opportunities and limitations of these tools. When presenting the case, the data scientist demonstrated the prototype (an AI system) under development, highlighting several of its outcomes. The prototype served as a boundary object for discussing current possibilities, limitations, data requirements, and potential improvements. To facilitate discussion of the AI system, the researchers had previously developed an origami toolkit with various paper tokens (see Figure 1A) to help participants explore the opportunity space for conceptualising improvements in manufacturing (Aranda-Muñoz et al., 2024). The operators used these paper tokens to describe the assembly process, the instruction software they use, and common assembly errors. Such representation through tokens enabled participants to ask questions about the context to understand it better and to share a common view of the opportunity with the AI system.

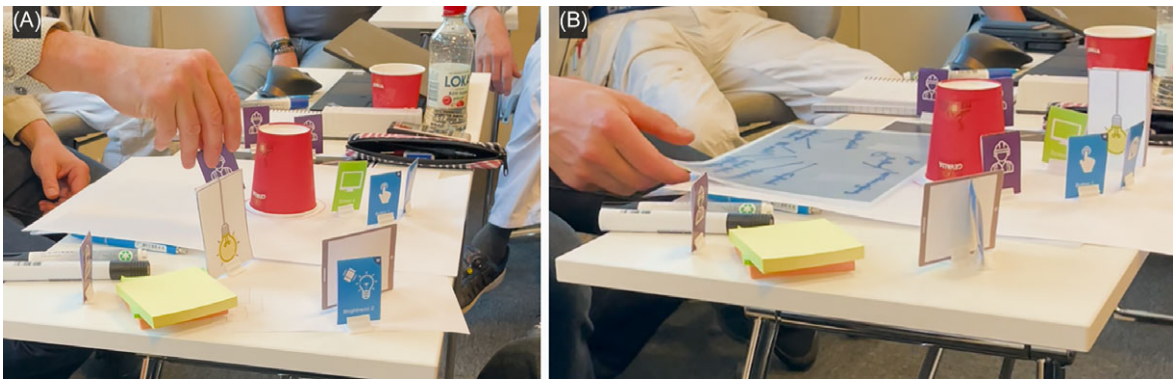


Figure 1. Participants explore the opportunity space (A) and annotate desirable AI outputs (B)

For the next activity, the interaction designer used photographs from the vision system (printed in colour on A4 paper) as boundary objects (Figure 1B). Participants annotated the photographs with plastic sleeves and coloured pens to indicate desirable system outputs, such as specific types of assembly errors to be detected. This hands-on exploration generated discussions of key topics: the information operators needed, common assembly errors, the timing of error detection, preferred notification methods, and how visual overlays could communicate problems. The emergence of these aspects prompted new conversations with other roles, namely data scientists, technicians, and participants in research and development, who reflected on feasibility and sought compromises to realise some of these aspects. This process bridged the gap between the visual reality of the shop floor and the computer vision requirements of the data scientists, creating a shared language for what constitutes an error. One example discussed in depth was determining when an assembly error is detected, as some errors can be difficult to identify early due to the nature of the assembly process.

Lastly, we presented screenshots of the operators' assembly instructions (obtained from the current system) to facilitate discussion of the integration of the new AI system. These images helped operators articulate their preferred interaction methods. They expressed a strong preference against adding another system or screen, which prompted discussions about whether and how the AI system's new capabilities could be incorporated into the existing assembly instruction program.

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

4.3. Workshop 3: early discussions on integrating the AI system

The third workshop aimed to facilitate discussions about integration aspects by applying selected Human-AI guidelines (PAIR, 2021) that the researchers considered relevant to the case. The three-hour workshop consisted of exploring and role-playing four scenarios previously prepared by the researchers. Rather than introducing the guidelines directly to participants, the research team developed scenarios addressing specific topics drawn from the guidelines. To ground the scenarios and support role-play, the interaction designer created high-fidelity mock-ups of the potential AI system, imitating the visual style of the existing assembly instruction software. These A3-printed interfaces reflected insights from the previous two workshops, assuming the AI system would be integrated into the assembly system.

Before the role-playing began, the interaction designer introduced the new interface. Participants were asked to imagine that they had noticed the interface had changed mid-assembly and to think aloud as they explored the mock-ups. The first mock-up displayed a modified version of their current assembly instruction system. Participants observed that the system now recognised the assembled components, identified which step operators were currently on (a new feature), and confirmed that all previous AI inspections were correct (also new). Operators could "click" on the current step to access additional AI system capabilities, revealing a mock-up showing an error-free assembly process. This interface contained several elements: a "live" image of the assembly process with a navigation bar for scrolling through the history of previous inspections; an "inspection passed" message with various KPIs from the system's checks (e.g., number, placement, positioning, width, and height of specific elements); and a raw image of the assembly process with the latest timestamp.

The interaction designer then facilitated four activities, each corresponding to a scenario depicting a different assembly mistake discussed in the second workshop. Each activity consisted of three parts: an introduction to the scenario (with corresponding interface mock-ups), a guided discussion with preplanned questions based on Human-AI guidelines, and a collaborative reflection in which participants could express additional ideas, concerns, and expectations.

By combining guided scenarios with interface mock-ups, participants discussed: roles involved in troubleshooting; troubleshooting timelines; responsibility for addressing assembly problems (e.g., whether operators would need to contact mentors or colleagues depending on severity); necessary actions to solve problems; information needed for decision-making; notification methods (e.g., sound, messages); handling false positives; acceptable AI error rates (to calibrate the reward function); how operators could provide feedback; consequences of the system failing to recognize mistakes; the impact of too many notifications and what would constitute "too many" in their context; alarm thresholds; and whether operators might trust the vision system excessively (e.g., leading to complacency) and what would constitute "too much" trust. The interface mock-ups tangibly reflected many of these aspects, displaying different assembly errors, KPIs, AI system confidence levels, true and false positives and negatives, operator messages, and interactions between instructions and the AI system.

Following the workshop, the interaction designer iterated the mock-up designs based on the feedback gathered. The resulting mock-ups captured the cross-functional team's decisions and were sent to the designated contact for additional feedback (particularly from end users) and adjustments. The mock-ups served as a first draft capturing early design decisions for the AI system development. The project leader (R&D) further documented the case to demonstrate its value to higher decision-makers. A few months later, the case was approved internally, and implementation is currently ongoing.

5. Results and learnings

To address the first research question, this section (and Table 1) summarises the boundary objects employed across the three workshops, describing how each boundary object supported stakeholder discussions and the design decisions that emerged. Our assessment of these outcomes is qualitative, evidenced by the tangible artefacts and decisions co-created during the workshops. In the first workshop (WK1), the swimlane diagram template and needs identification template helped participants understand the AI system's opportunity space and identify key stakeholders. The second workshop (WK2) employed more concrete artefacts: an origami toolkit, photographs of assembly stages, screenshots of assembly instructions (all of which helped lower the barrier to participation and build shared understanding of the operators' actual work context), as well as Teachable Machine and the AI system prototype to demonstrate capabilities and expose the limitations of AI vision technology. Because AI was demystified as an imperfect tool, in the third workshop (WK3), the operators could discuss with other

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

roles on acceptable AI prediction levels, confidence thresholds, false-positive/negative trade-offs, and feedback workflows, resulting in specific design decisions, such as colour-coded confidence levels and protocols for operator feedback. Such discussions were facilitated by high-fidelity UI mock-ups and four scenarios that incorporated these design aspects. A telling example is that the UI confidence indicators conveyed trust and system reliance to operators, while simultaneously representing statistical thresholds and false-positive rates to data scientists, thereby allowing the cross-functional team to negotiate shared work arrangements, such as the feedback workflow, despite differing technical conceptualisations of the system. The emergence of a shared language for defining assembly errors, alongside cross-functional agreement on colour-coded confidence thresholds and feedback workflows, provides empirical evidence that a shared understanding was successfully negotiated among the different roles. Across all workshops, boundary objects thus moved from abstract exploration to concrete design decisions, enabling the cross-functional team to externalise and negotiate their needs, constraints, and expectations for the AI system development and integration.

Table 1. Overview of Boundary Objects and their observed support and outcomes

Boundary Object	Observed Support	Observed Outcomes
WK1: Swimlane Diagram Template	Facilitated shared understanding of the AI system's opportunity space and identification of affected roles.	Key roles identified; first draft of workflow showing system integration with roles and processes.
WK1: Template for Needs Identification & Prioritization	Enabled early discussions on opportunities considering stakeholder mental models, AI prototype, training data, and user interface.	Documented needs (especially for end-users and interface features); prioritized action list for next workshop.
WK2: Generic Demonstrator, Teachable Machine	Demystified AI vision and prompted discussions on AI limitations and model evolution.	Participants understood vision system capabilities and limitations; recognized need for operator feedback and cross-role collaboration.
WK2: AI System Prototype	Introduced opportunity to support operators in detecting assembly mistakes; demonstrated current features and training data.	Operators onboarded to prototype; shared understanding of opportunity space and challenges (e.g., recognition of limitations, data collection).
WK2: Origami Business Toolkit	Enabled operators to present their assembly process using paper tokens and explore integration alternatives.	Collective shared understanding of context and potential integration into assembly line.
WK2: Photographs of Assembly Stages + Plastic Sleeves	Supported discussions on operator needs, notifications, and data collection for specific mistakes.	Seven additional features identified; operators understood need to provide feedback and photographs.
WK2: Screenshots of Assembly Instructions	Explored integration alternatives for operators' work practices.	Decision to integrate AI recommendations into existing instruction software.
WK3: UI Mock-ups + Four Scenarios	Facilitated discussions on acceptable AI prediction levels, confidence levels, thresholds, balancing true and false positives. Addressed challenges of operators providing feedback given assembly problem resolution timelines.	Awareness of need to balance false alarms and missed mistakes; shared understanding of collaborative approach needed. Decision to display confidence levels using color-coded thresholds. Early workflow for feedback process; decisions on who/when reports feedback.

To address the second research question, the boundary objects supported different stakeholder roles in different ways (see Table 2). For operators, the boundary objects served multiple functions: demystifying AI vision technology and exploring its limitations through a demonstrator; onboarding them to the AI system concept through a prototype; managing expectations about what the AI system could achieve; helping them understand their role during development (e.g., articulating their needs) and after deployment (e.g., providing feedback); and supporting their externalization of contextual aspects of the assembly process through the origami toolkit. Several operators and R&D participants noted that the

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

mock-up interfaces and scenarios grounded discussions of potential system features, even when those features had not yet been implemented in the prototype. These discussions enabled operators to articulate contextual requirements for integrating the solution into the assembly line and to clarify workflows among operators, quality technicians, and R&D participants.

The data scientist gained a deeper understanding of contextual aspects and operator needs, established initial collaborations with other team members for follow-up meetings, and learned how collaborative methods and materials helped externalise aspects for discussion and reach shared understandings.

The project leader received support to align the cross-functional team and obtained tangible materials to externalise and document requirements, AI model feasibility, and potential integration into the assembly line. Some activities and materials (e.g., swimlanes and origami toolkit) helped identify and onboard key stakeholders who could be affected by the integration. Other benefits included gaining nuanced contextual insights to develop a more concrete pitch to decision-makers to secure approval for the next implementation phase. Moreover, the project leader and data scientist emphasised in their reflections the importance of maintaining this collaborative, cross-functional approach to address assembly-line needs while ensuring that AI system features remain feasible.

Table 2. Observations and reflections on how boundary objects helped the cross-functional team

Role	Researchers' observations and reflections from participants using boundary objects
Operators (end-users)	First contact point with AI vision and demystification of what the technology does. Onboarding process of the potential AI system prototype: initial mental model. Management of expectations of what is possible from the AI model. Understanding they can influence the desired features of the system, which is their role in the development and maintenance of the system (e.g., feedback), and how to collaborate with development (e.g., document assembly errors). Support to explain and externalize contextual aspects of their work activities and context.
Data Scientist	Learning a collaborative method to engage end-users in the development. Spread the understanding of the potential AI system. Deeper understanding of what features would be more valuable for the operators. Contextual understanding of how the AI system can be integrated into the assembly line. Collaboration with research and development and operators to gather more data. Collaboration with the operators (end-users) to discuss further on topics (e.g., confidence levels, thresholds, KPIs, desired features, etc.)
Project Leader (R&D)	Help in orchestrating the requirements, feasibility and potential integration of the ML model. Help to identify and onboard relevant stakeholders into the early design and development process of the AI system. Help to create a case to "pitch" for higher decision-makers to start the implementation stage. Shared view: all stakeholders need to contribute with knowledge.
Head of assembly	Insights into the workflows and who needs to be involved when alarms happen. Insights into how the potential AI system affects the responsibility of operators. Help to assess the potential cost savings if the AI system would be integrated. Help to reflect on some of the limitations and potential challenges of the AI system.
Research & development, technicians.	Onboarding to the AI system and needs from different roles. Understand how to help operators and data scientist in the identification of assembly mistakes for data needs of the AI system.

The head of assembly benefited from insights generated through explorations with boundary objects, particularly regarding how the new workflow would function when the AI system is integrated, considering each role's responsibilities and potential impacts (e.g., through swimlanes, origami toolkit, mock-ups, and scenarios). This information helped the cross-functional team better assess potential savings from integrating the AI system while identifying implementation challenges (e.g., capturing photographs during assembly) and post-integration issues (e.g., reward function calibration).

Lastly, R&D roles and technicians benefited from onboarding to the AI system and from understanding the distinct needs of different roles within the cross-functional team. This onboarding clarified their roles during development (e.g., identifying and categorising assembly errors) and post-integration (e.g., supporting calibration and responding to alarms).

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

6. Discussion

This study contributes to the literature on collaborative methods for aligning user needs with AI systems by documenting how Human-AI guidelines can be embedded in boundary objects that facilitate cross-functional collaboration during early AI system design and development in manufacturing. Drawing on a real-world AI case, the study underscores the practical utility of specific boundary objects (including AI demonstrators, prototypes, an origami toolkit, annotated photographs, and high-fidelity mock-ups with role-play scenarios) across three workshops with diverse industrial stakeholders. While previous research has established the value of boundary objects in AI development in domains such as healthcare, this study suggests manufacturing can adopt and adapt these insights, with integration into the factory floor and alignment with work-oriented practices as important considerations. More particularly, this case documents how a cross-functional team used boundary objects to collaboratively explore the opportunity space of an AI system idea, attending to the needs of different stakeholders, including end-user, AI model, data, and workflow requirements, alongside broader integration considerations. In doing so, it reflects on how each role in the team member benefited from the workshop materials and activities, thereby addressing both research questions. This section briefly outlines broader reflections, lessons learned, and future directions.

From an HCI research perspective, we identify a need for more systematic approaches to understanding how boundary objects support cross-functional teams during the early design and development stages of AI systems. For example, we see the need for theoretical frameworks that describe the roles that boundary objects play at different stages of design and the roles for which they serve. Examples of some roles of boundary objects that we observed in the case study include demystifying AI capabilities for non-experts (e.g., operators could experience firsthand prototypes to understand data limitations and their role in the feedback loop), embedding discussions on particular Human-AI guidelines, experiencing and assessing AI models, exploring what data is necessary, and sharing common views along the different design stages, among others; ideally, such boundary objects could be mapped into different stages of design (e.g., double diamond) to help future design practitioners support the journey from AI ideation to integration.

From a broader perspective, stepping back from boundary objects, we observe a lack of methodologies and frameworks to systematically address the early ideation and development stages of AI systems. New literature addressing these stages is promising (Yildirim, 2024; Yildirim et al., 2024; Yildirim, Oh, et al., 2023; Yildirim, Pushkarna, et al., 2023) but remains scarce. More studies are needed that document these early stages to identify which contextual aspects emerge from the use of Human-AI guidelines, with examples of how HCI researchers can adapt such guidelines to different domains and describe which are useful and which are discarded based on the particulars of the context.

Manufacturing AI development frameworks (Yamamoto et al., 2025) share some similarities with those in healthcare, yet a considerable gap remains between the two domains. Healthcare research has produced holistic AI development frameworks (Assadi et al., 2022; Sande et al., 2022), examples of scaffolding ideation (Yildirim et al., 2023), human-AI onboarding materials (Cai et al., 2019, 2021), approaches to identifying and designing vision system applications (Yildirim et al., 2024), and insights into how such systems can be integrated into practice (Yang, Steinfeld, et al., 2019). Based on our review of AI development frameworks in manufacturing (Yamamoto et al., 2025) and previous challenges of adopting HAI guidelines (Yamamoto et al., 2024), these aspects of the early design stages of AI development remain largely undocumented in the manufacturing sector.

Lastly, we would like to underscore the opportunity for designers to orchestrate and facilitate cross-functional teams during the early design and development stages of AI systems. Recent literature has begun to recognise the opportunity (Yildirim, 2024, p. 41), with discussions on the knowledge required of designers to lead the development of AI solutions (Colombo & Costa, 2024). This study underscores the role of the interaction designer, who designed several boundary objects and facilitated several workshops. Such workshops can be the starting point for establishing the cross-functional team and approaching AI system development from a collaborative perspective between different stakeholders, where designers can materialize and embed Human-AI guidelines into boundary objects to support necessary dialogues for capturing needs, understanding necessary collaborations between different job roles, address and mitigate potential challenges of AI systems, and maintain a dialogue between the AI system feasibility aspects and human and collaborative perspectives. These dialogues are necessary for AI innovation and have been described as gaps

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

between technology-centric and user-centric approaches in AI (Yang, 2021, p. 13; Yildirim, 2024, p. 5), in which designers are considered vital in the ideation and problem formulation stages (Yildirim, 2024, p. 5). Future research can investigate the skills designers need to acquire and the other roles that can lead and support the orchestration of the different stages of AI development and integration within cross-functional teams.

7. Conclusion

This case study documented the early design and development of an AI vision system for an assembly line at a manufacturing company producing electrical equipment. Our findings indicate that boundary objects were beneficial to the cross-functional team, as they facilitated shared understanding of the potential AI system, helped explore and define end-user needs, discussed potential AI system challenges and planned ways to overcome them, and established collaborations among key roles for follow-up activities. The case study underscores the need for close collaboration in the cross-functional team in the early design and development stages of AI systems. Such close collaboration can be supported by designers, who can create boundary objects and orchestrate shared explorations and understandings during the initial stages of AI design and development.

Acknowledgement

PiiA and Produktion2030 funded this research under the project Digitala Stambanan.

References

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for Human-AI Interaction. *Proceedings of 2019 CHI*. https://doi.org/10.1145/3290605.3300233

Aranda-Muñoz, A., Yamamoto, Y., & Sandström, K. (2024). The Karakuri IoT toolkit: a collaborative solution for ideating and prototyping IoT opportunities. *Design 2024, 18th International Design Conference, Croatia.* https://doi.org/10.1017/pds.2024.21

Assadi, A., Laussen, P. C., Goodwin, A. J., Goodfellow, S., Dixon, W., Greer, R. W., Jegatheeswaran, A., Singh, D., McCradden, M., Gallant, S. N., Goldenberg, A., Eytan, D., & Mazwi, M. L. (2022). An integration engineering framework for machine learning in healthcare. *Frontiers in Digital Health*, 4. https://doi.org/10.3389/fdgth.2022.932411

Ayobi, A., Stawarz, K., Katz, D., Marshall, P., Yamagata, T., Santos-Rodríguez, R., Flach, P., & O'Kane, A. A. (2021). Machine Learning Explanations as Boundary Objects: How AI Researchers Explain and Non-Experts Perceive Machine Learning. *2021 Joint ACM Conference ACMIUI-WS 2021*. http://ceur-ws.org/Vol-2903/IUI21WS-TExSS-3.pdf

Cai, C. J., Winter, S., Steiner, D., Wilcox, L., & Terry, M. (2019). "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proc. ACM Human-Computer Interaction, 3(CSCW).* DOI: https://doi.org/10.1145/3359206

Cai, C. J., Winter, S., Steiner, D., Wilcox, L., & Terry, M. (2021). Onboarding Materials as Cross-functional Boundary Objects for Developing AI Assistants. *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3411763.3443435

Carney, M., Webster, B., Alvarado, I., Phillips, K., Howell, N., Griffith, J., Jongejan, J., Pitaru, A., & Chen, A. (2020). Teachable Machine: Approachable Web-Based Tool for Exploring Machine Learning Classification. *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-8. https://doi.org/10.1145/3334480.3382839

Colombo, S., & Costa, C. (2024). Can designers take the driver's seat? A new human-centered process to design with data and machine learning. *The Design Journal*, 27(1), 7–29. https://doi.org/10.1080/14606925.2023.2279835

Haefner, N., Parida, V., Gassmann, O., & Wincent, J. (2023). Implementing and scaling artificial intelligence: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 197, 122878. https://doi.org/10.1016/j.techfore.2023.122878

Jansen, A., & Colombo, S. (2023). Mix & Match Machine Learning: An Ideation Toolkit to Design Machine Learning-Enabled Solutions. *Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction.* https://doi.org/10.1145/3569009.3572739

Makarius, E. E., Mukherjee, D., Fox, J. D., & Fox, A. K. (2020). Rising with the machines: A sociotechnical framework for bringing artificial intelligence into the organization. *Journal of Business Research*, 120, 262–273. https://doi.org/10.1016/j.jbusres.2020.07.045

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press

Nahar, N., Zhou, S., Lewis, G., & Kästner, C. (2022). Collaboration challenges in building ML-enabled systems: communication, documentation, engineering, and process. *Proceedings of the 44th ICSE*, 413-425. https://doi.org/10.1145/3510003.3510209

PAIR, G. (2021). *People + AI Guidebook*. URL: https://pair.withgoogle.com/guidebook

Sande, D. van de, Genderen, M. E. V., Smit, J. M., Huiskens, J., Visser, J. J., Veen, R. E. R., Unen, E. van, BA, O. H., Gommers, D., & Bommel, J. van. (2022). Developing, implementing and governing artificial intelligence in medicine: a step-by-step approach to prevent an artificial intelligence winter. *BMJ Health & Care Informatics*, 29(1), e100495. https://doi.org/10.1136/bmjhci-2021-100495

Sendak, M. P., Ratliff W., Sarro, D., Alderton, E., Futoma, J., Gao, M., Nichols, M., Revoir, M., Yashar, F., Miller, C., Kester, K., Sandhu, S., Corey, K., Brajer, N., Tan, C., Lin, A., Brown, T., Engelbosch, S., Anstrom, K., Elish, M. C., Heller, K., Donohoe, R., Theiling, J., Poon, E., Balu, S., Bedoya, A., O'Brien, C. (2020). Real-World Integration of a Sepsis Deep Learning Technology Into Routine Clinical Care: Implementation Study. *JMIR Med Inform*, 8(7), e15182. https://doi.org/10.2196/15182

Star, S. L., & Griesemer, J. R. (1989). Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, 19(3), 387–420. https://doi.org/10.1177/030631289019003001

Subramonyam, H., Im, J., Seifert, C., & Adar, E. (2022). Solving Separation-of-Concerns Problems in Collaborative Design of Human-AI Systems through Leaky Abstractions. *Proceedings of the 2022 CHI*. https://doi.org/10.1145/3491102.3517537

Yamamoto, Y., Aranda-Muñoz, Á., & Sandström, K. (2025). A development framework for human work integrated AI systems in manufacturing. *Frontiers in Manufacturing Technology*. https://doi.org/10.3389/fmtec.2025.1601903

Yamamoto, Y., Muñoz, Á. A., & Sandström, K. (2024). Challenges in designing a Human-centred AI system in manufacturing. *International Journal of Mechatronics and Manufacturing Systems*. https://doi.org/10.1504/IJMMS.2024.144289

Yang, Q. (2021). *Profiling Artificial Intelligence as a Material for User Experience Design* [Doctoral thesis, Carnegie Mellon University]. https://doi.org/10.1184/R1/14376731

Yang, Q., Cranshaw, J., Amershi, S., Iqbal, S. T., & Teevan, J. (2019). Sketching NLP: A Case Study of Exploring the Right Things To Design with Language Intelligence. *Proceedings of the 2019 CHI*, 1-12. https://doi.org/10.1145/3290605.3300415

Yang, Q., Steinfeld, A., & Zimmerman, J. (2019). Unremarkable AI: Fitting Intelligent Decision Support into Critical, Clinical Decision-Making Processes. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-11. https://doi.org/10.1145/3290605.3300468

Yildirim, N. (2024). *Discovering the Right Things to Design with Artificial Intelligence* [Doctoral thesis, Carnegie Mellon University]. https://doi.org/10.1184/R1/26485180

Yildirim, N., Oh, C., Sayar, D., Brand, K., Challa, S., Turri, V., Walton, N. C., Wong, A. E., Forlizzi, J., McCann, J., & Zimmerman, J. (2023). Creating Design Resources to Scaffold the Ideation of AI Concepts. *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, 2326-2346. https://doi.org/10.1145/3563657.3596058

Yildirim, N., Pushkarna, M., Goyal, N., Wattenberg, M., & Viégas, F. (2023). Investigating How Practitioners Use Human-AI Guidelines: A Case Study on the People + AI Guidebook. *Proceedings of the 2023 CHI*. https://doi.org/10.1145/3544548.3580900

Yildirim, N., Richardson, H., Wetscherek, M. T., Bajwa, J., Jacob, J., Pinnock, M. A., Harris, S., Castro, D. C. D., Bannur, S., Hyland, S., Ghosh, P., Ranjit, M., Bouzid, K., Schwaighofer, A., Pérez-García, F., Sharma, H., Oktay, O., Lungren, M., Alvarez-Valle, J., … Thieme, A. (2024). Multimodal Healthcare AI: Identifying and Designing Clinically Relevant Vision-Language Applications for Radiology. *Proceedings of the CHI Conference*. https://doi.org/10.1145/3613904.3642013

https://doi.org/10.1017/pds.2026.10373 Published online by Cambridge University Press