This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

Scaling AI Transformation in Large Organisations through Peer Learning: A Case Study at Deutsche Telekom

Stefan Kohn *

Hochschule Darmstadt, Max-Planck-Str. 2, 64807 Dieburg, Germany.
European University of Technology, European Union.
E-mail: Stefan.kohn@h-da.de

Stefan Dörken

Deutsche Telekom Service GmbH, Friedrich-Ebert-Allee 71-77, 53113 Bonn, Germany.
E-mail: stefan.doerken@telekom.de

Sindy Leffler-Krebs

Deutsche Telekom AG, Friedrich-Ebert-Allee 140, 53113 Bonn, Germany.
E-mail: Sindy.Leffler-Krebs@telekom.de

Arno Selhorst

Deutsche Telekom Service GmbH, Friedrich-Ebert-Allee 71-77, 53113 Bonn, Germany.
E-mail: Arno.Selhorst@telekom.de

* Corresponding author

Abstract: Large organisations face a fundamental challenge in scaling artificial intelligence (AI) transformation beyond isolated expert initiatives and episodic learning formats. While generative AI tools are increasingly accessible, their effective and responsible integration into everyday work requires broad-based capability building, contextual application knowledge, and sustained organisational learning. This paper presents a qualitative practitioner case study of Deutsche Telekom's AI Pioneers program, a peer-learning initiative designed to scale AI literacy and application capabilities across a large organisation. Drawing on semi-structured expert interviews, internal program documentation, and descriptive participant feedback, the paper analyses how peer learning enables decentralised knowledge creation, cross-functional diffusion, action-oriented experimentation, and role-based scaling. The findings show that peer learning can serve as a powerful mechanism for AI transformation when embedded in real work contexts. At the same time, the case demonstrates that decentralised learning does not scale autonomously. Active community management is required to sustain engagement, provide orientation, and translate learning into organisational impact.

Keywords: AI Transformation; Peer Learning; AI Literacy; Case Study; Organisational Learning; Experiential Learning; Community Management; Capability Building; Deutsche Telekom; Responsible AI.

1 Introduction

The rapid development of artificial intelligence, and especially generative AI, has shifted the innovation agenda of large organisations. AI is no longer confined to expert systems, data science departments, or specialised automation projects. Instead, AI tools increasingly become accessible to broad groups of employees and can be applied in everyday tasks such as writing, summarizing, searching, translating, ideating, analysing, coding, documenting, and supporting decision-making. This democratization of AI creates a new organisational challenge. The question is no longer only how to develop AI systems, but how to enable a large and heterogeneous workforce to use AI effectively, responsibly, and productively.

This challenge is particularly acute in large organisations. Their size creates both the need and the difficulty of scaling AI capabilities. On the one hand, large organisations possess diverse operational contexts in which AI can create value. On the other hand, these same organisations are characterised by functional specialisation, hierarchical structures, varying digital maturity, and different levels of openness toward technological change. As a result, AI transformation often remains fragmented. Some teams experiment intensively, while others remain uncertain about what AI means for their work. Central expert teams can provide guidance, tools, and governance, but they cannot identify, evaluate, and implement all relevant use cases across every function.

The literature on AI literacy increasingly emphasizes that broad AI capability building cannot be reduced to technical training. AI literacy includes the ability to understand AI's potential and limitations, apply AI tools in concrete contexts, critically evaluate AI-generated outputs, and use AI responsibly (Long and Magerko, 2020; Ng et al., 2021). Recent literature also points toward participatory, experiential, and context-embedded learning approaches as particularly promising for organisational AI literacy. Existing studies converge on the importance of employee participation, hands-on experimentation, and task relevance, but evidence on enterprise-wide implementation at scale remains limited. Most documented approaches remain small-scale pilots, department-level interventions, or conceptual frameworks rather than large-scale organisational programs (Rueckert et al., 2021; Ruess et al., 2024; Schoenmaker, 2021).

Deutsche Telekom's earlier Promptathon initiative addressed part of this challenge by creating a low-threshold, gamified format for hands-on AI learning. Promptathons helped employees experience generative AI directly, develop confidence, and connect AI literacy with responsible use. Since 2023, this format has engaged thousands of employees and demonstrated that experiential learning can be effective at scale when designed as an accessible, social, and practice-oriented intervention (Kohn and Leffler-Krebs, 2025). However, Promptathons are event-based. They can trigger awareness, motivation, and initial competence, but they do not automatically create sustained AI capability within teams and business units.

This paper therefore examines the next stage of AI capability building at Deutsche Telekom: the AI Pioneers program. The program aims to move beyond episodic learning events toward a more continuous and decentralised model of AI transformation. It is built around the idea that AI knowledge should not remain concentrated in central expert teams, but should be distributed through peer learning, community structures, and locally embedded multipliers.

The guiding research question is: How can peer learning enable scalable AI transformation in large organisations?

The contribution of this paper is practical and conceptual. Practically, it describes how a large organisation operationalised peer learning for AI transformation through a structured multi-level learning journey and community-based approach. Conceptually, it identifies mechanisms through which peer learning supports AI transformation and shows why peer learning requires active orchestration through community management. The central insight is that peer learning can scale AI transformation, but it does not scale by itself.

2 Background: AI literacy, experiential learning, and the scaling gap

AI transformation depends on more than the availability of AI tools. It requires employees to develop new forms of competence. Long and Magerko define AI literacy as a set of competencies that enables individuals to critically evaluate AI technologies, communicate and collaborate with AI, and use AI as a tool in everyday contexts (Long and Magerko, 2020). Ng et al. further distinguish several components of AI literacy, including understanding AI concepts, using and applying AI, evaluating and creating with AI, and ethical awareness (Ng et al., 2021). These definitions are useful because they frame AI literacy not as narrow technical expertise, but as a broader capability that combines knowledge, application, judgement, and responsibility.

For large organisations, this broader understanding is essential. Most employees do not need to become AI developers, but many need to understand how AI can support their work, where its limitations lie, and how to use it responsibly within organisational and regulatory boundaries. In this sense, AI literacy is both an individual and an organisational capability. It enables employees to act, but it also determines whether an organisation can move from isolated AI experiments to systematic adoption.

The reviewed literature points toward several principles for building AI literacy in organisations. Participatory design appears as a recurring theme. Employees should not be passive recipients of finished systems or abstract training content; they should be involved in identifying needs, shaping applications, and reflecting on consequences. Experiential learning is similarly emphasized. Employees learn more effectively when they can experiment with AI systems, develop prototypes, or apply AI to concrete work problems. Social and interprofessional learning also appear as important, especially where AI use requires collaboration between domain experts, technical experts, and organisational stakeholders (Antweiler et al., 2023; Pfeiffer, 2020; Rueckert et al., 2021; Ruess et al., 2024; Schoenmaker, 2021).

Several examples illustrate these principles. Rueckert et al. describe the KI_Cafe method as an on-site learning factory that allows employees at different organisational levels to engage with AI through examples, exhibits, and co-development of AI-based assistance systems (Rueckert et al., 2021). Ruess et al. develop a participatory model for AI implementation in business processes, emphasizing early involvement, qualification, and low-barrier formats connected to employees' actual tasks (Ruess et al., 2024). Pfeiffer highlights the importance of employees' existing context competence and argues that organisations often underestimate the AI potential already present in their workforce (Pfeiffer, 2020). Antweiler et al. show, in the healthcare context, that AI-related capability building must include digital literacy, interdisciplinary collaboration, and attention to role changes and risks (Antweiler et al., 2023).

This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

At the same time, the literature reveals a significant gap. While participatory and experiential approaches are conceptually well supported, rigorous outcome evidence remains limited. There are few robust pre-post comparisons, controlled evaluations, or long-term follow-up studies on AI literacy gains. Moreover, many implementations described in the literature are not enterprise-wide. They often focus on SMEs, specific departments, or professional domains. Large-scale, company-wide initiatives such as Promptathons or peer-learning programs in large corporations remain underrepresented.

This gap matters for innovation management. Scaling AI transformation requires more than designing an effective workshop. It requires a learning architecture that can operate across organisational boundaries, remain relevant for different work contexts, and sustain engagement beyond initial enthusiasm. Peer learning offers one possible answer because it mobilizes employees as both learners and knowledge carriers. Yet peer learning itself raises a further question: how can it be structured, supported, and sustained in a large organisation?

3 Methodology

This paper adopts a qualitative practitioner case study approach. The aim is not to test hypotheses or produce statistically generalizable results, but to generate practice-based insights into how peer learning can support AI transformation in a large organisational context. This positioning is important because the available empirical material is rich in qualitative and descriptive evidence, but it does not constitute a quantitative impact study.

The case is Deutsche Telekom's AI Pioneers program. The program was selected because it explicitly addresses the challenge of scaling AI capabilities through decentralised learning, peer exchange, and community-based diffusion. It also represents a follow-up development to Deutsche Telekom's Promptathon initiative, which focused on broad initial AI literacy through hands-on events (Kohn and Leffler-Krebs, 2025).

The study draws on three sources of evidence. First, three semi-structured expert interviews were conducted with key stakeholders involved in the design, implementation, and development of the AI Pioneers program. The interview guide covered the origin of the initiative, program structure, learning formats, knowledge diffusion, success factors, challenges, leadership support, impact, and future development. Second, internal program documents were analysed to understand the structure, objectives, learning journey, participation levels, and development logic of the program. Third, internal participant feedback and program indicators were used descriptively to support the qualitative interpretation. These indicators include participation data, completion numbers, mean recommendation ratings on a ten-point scale, and self-reported usefulness ratings on five-point scales.

The survey data is used cautiously. The data does not include pre-post measurements, control groups, multi-item scales, or statistical tests. Its function is therefore not to prove causal impact, but to triangulate and illustrate patterns that emerge from interviews and program documentation.

The contribution lies in analysing a real-world implementation and deriving design principles for practice. The paper therefore follows the logic of an innovation-in-practice case: it describes an organisational issue, explains how the organisation responded, analyses what was learned, and derives implications for others facing similar transformation challenges.

4 Case description: The AI Pioneers program

The AI Pioneers program emerged from a practical scaling problem. Before the program was launched, AI enablement at Deutsche Telekom was already taking place through presentations, workshops, Promptathons, and communication activities. These formats created awareness and generated demand. However, they also revealed the limits of centralised enablement. Employees repeatedly asked what AI could mean for their specific area, process, team, or function. A small group of central experts could answer general questions about AI, governance, tools, or prompting, but it could not identify and develop use cases for every business area.

The core insight was that AI application knowledge is distributed. Finance, service, HR, network operations, compliance, marketing, and project management each have their own tasks, constraints, data environments, workflows, and value-creation logics. Relevant AI use cases therefore cannot be fully prescribed from the centre. They need to be discovered and shaped where the domain knowledge exists. The AI Pioneers program was created to activate this distributed knowledge and combine it with AI literacy.

The program is structured as a multi-level learning journey open to all employees of Deutsche Telekom. Through this approach, the program was designed to spread AI literacy across the organisation and reach as many teams as possible. Level 1, Explorers, focuses on basic orientation, responsible AI use, prompting, exchange, AI updates, and individual learning pathways. Level 2, Architects, focuses on identifying use cases, evaluating tools, moderating AI-related discussions, recognizing pain points, creating opportunities, and moving from ideas toward implementation. Level 3, Visionaries, focuses on strategic perspective, maturity assessment, decision logic, and project implementation. The levels are designed to move participants from initial understanding toward team enablement and project-oriented transformation.

Table 1 Structure of the AI Pioneers learning journey

Level	Focus	Intended transformation role
Level 1: Explorers	Orientation, responsible AI, prompting, exchange, AI updates	Build initial AI literacy and lower barriers to use
Level 2: Architects	Use-case identification, tool evaluation, moderation, structured implementation	Enable participants to bring AI competence into their teams
Level 3: Visionaries	Strategic perspective, maturity assessment, project implementation	Connect AI learning with organisational and project impact

Source: Authors' summary based on internal program documentation.

The program reached substantial scale in its first phase. In 2025, 718 employees were trained in Level 1 and 156 in Level 2. The first Level 3 cohort trained 70 participants in 2026. The overall internal documentation reports multiple cohorts and a visible progression from broad entry formats toward smaller, more advanced levels. The average recommendation ratings were high across levels: 8.58 for Level 1, 8.61 for Level 2, and 8.27 for Level 3 on a ten-point scale. Level 1 feedback also showed positive self-reported

This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

ratings for perceived support in handling AI and usefulness for work, generally around four on a five-point scale. These values should not be overinterpreted as impact measures, but they indicate that participants perceived the program as relevant and useful.

Table 2 Descriptive participation and feedback indicators

Program level	Participation / completion	Mean recommendation rating
Level 1: Explorers	718 participants trained in 2025	8.58 on a 10-point scale
Level 2: Architects	156 participants trained in 2025	8.61 on a 10-point scale
Level 3: Visionaries	70 participants trained in 2026	8.27 on a 10-point scale

Source: Internal program documentation, 2025/2026.

The learning design combines structured modules with peer exchange, group work, practical exercises, and community elements. Across the feedback documents, participants repeatedly mention exchange, networking, hands-on work, real use cases, practical relevance, and learning from other areas as central benefits. At the same time, they also request more time, clearer structure, better expectation management, more concrete examples, and stronger support for translating ideas into implementation.

This combination of strong perceived value and recurring requests for more structure is central to the case. It shows that peer learning creates energy and relevance, but it also creates coordination needs.

5 Findings: How peer learning enables scalable AI transformation

The AI Pioneers case shows that peer learning is not merely a didactic technique. It functions as an organisational scaling mechanism. Through the program, AI knowledge is not simply transmitted from experts to employees; it is created, interpreted, adapted, and circulated among employees. Five mechanisms explain how this contributes to scalable AI transformation.

The first mechanism is localized knowledge creation. AI use cases are most meaningful when they emerge from concrete work contexts. Central teams can explain general AI principles, provide approved tools, and define governance requirements, but they usually lack the detailed operational knowledge required to identify high-value use cases in every business area. The AI Pioneers program addresses this by shifting part of the discovery process to employees themselves. Participants are not only trained in AI; they are encouraged to connect AI with their own workflows, pain points, and team challenges.

This mechanism is closely aligned with the distinction between AI/ML competence and context competence described by Pfeiffer. Employees often possess deep knowledge of processes, customers, constraints, and informal routines. AI transformation becomes more realistic when this context competence is treated as a resource rather than as an implementation obstacle (Pfeiffer, 2020). In the AI Pioneers program, the relevant question is not only what AI can do, but what AI can do here, in this task, with these constraints, for this team. This shift is essential for moving AI from abstract potential to organisational relevance.

The second mechanism is peer-driven knowledge diffusion. The internal feedback across the program repeatedly identifies exchange with colleagues as one of the most valuable elements. Participants benefit from hearing how other teams interpret AI, which tools they use, what obstacles they face, and which approaches have worked in practice. This type of learning is difficult to achieve through standardized training materials because much of the relevant knowledge is tacit, emerging, and context-dependent.

Peer diffusion is particularly important in large organisations because it allows knowledge to travel laterally. Instead of relying only on hierarchical communication or central training campaigns, knowledge spreads through relationships, communities, and shared activities. Participants learn not only from instructors but from each other. This creates a more flexible and adaptive knowledge system. It also supports cross-functional learning: employees discover that similar problems exist in different units and that solutions developed in one area may inspire applications elsewhere.

The third mechanism is action-oriented learning. The program's strongest feedback themes are not abstract knowledge gains but practical relevance, exercises, use-case work, and direct applicability. This reflects a broader pattern in the literature: AI literacy is more likely to become meaningful when employees can try, test, observe, and reflect rather than only listen. Rueckert et al.'s KI_Cafe approach, Schoenmaker's prototype-based field trials, and the Promptathon case all point in the same direction: learning by doing is particularly powerful in AI contexts because AI capabilities and limitations become visible through interaction (Rueckert et al., 2021; Schoenmaker, 2021; Kohn and Leffler-Krebs, 2025).

In the AI Pioneers program, action-oriented learning also serves another function. It lowers the psychological threshold for AI use. Many employees approach AI with uncertainty, scepticism, or inflated expectations. By working with concrete tasks and examples, participants can develop a more realistic understanding. They experience both usefulness and limitations. This supports responsible adoption because it counters two equally problematic patterns: fear-driven rejection and uncritical hype.

The fourth mechanism is role-based scaling. The term AI Pioneer signals that participants are not only learners, but potential multipliers. The program's ambition is not fulfilled when one employee completes a module. Its transformational potential depends on whether participants bring knowledge into their teams, initiate conversations, support colleagues, identify use cases, and model responsible AI behaviour. This multiplication logic responds directly to the limits of central enablement. A small expert group cannot accompany every team, but a distributed network of pioneers can make AI conversations more local, credible, and continuous.

However, this mechanism is not automatic. Interviews and feedback indicate that participants differ substantially in motivation, confidence, time availability, and willingness to take initiative. Some participants actively support their teams; others consume the program more passively. This variation is crucial. It shows that assigning a multiplier role does not ensure multiplier behaviour. Role-based scaling requires activation, support, expectations, and recognition.

The fifth mechanism is community-based learning. The AI Pioneers program creates a cross-functional social space in which AI-interested employees can connect. The community provides access to expertise, creates motivation, and makes learning visible. It also offers a structure for continuing exchange beyond individual modules. The internal

This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

program documentation explicitly emphasizes community spirit, knowledge sharing, responsible AI, innovation, curiosity, and the ambition to make AI accessible to many rather than few.

This community dimension is important because AI transformation is a moving target. Tools, regulations, internal policies, and use cases change rapidly. A one-time training can become outdated quickly. A community, by contrast, can adapt. It allows members to share updates, interpret new developments, and collectively make sense of changing possibilities. In this sense, the community is not an add-on to the learning journey. It is part of the infrastructure through which AI transformation can remain current.

6 Challenges: The limits of decentralised peer learning

The case also shows that peer learning has limits. The AI Pioneers program is not a simple success story in which community dynamics automatically solve the scaling problem. It reveals several tensions that are likely to occur in other large organisations as well.

The first tension concerns participation. The program is voluntary and attracts many motivated employees, but not all participants behave as active pioneers. Some enter with expectations shaped by traditional training: they expect content to be delivered, instructions to be clear, and next steps to be defined by others. This can conflict with the program's ambition to foster self-directed exploration and peer-driven contribution. Interview material highlights this challenge directly: part of the difficulty lies in moving employees from a passive learning mode toward active ownership.

This is not a weakness of the participants; it is an organisational reality. Large organisations have trained employees over many years to operate within established structures, roles, approval processes, and responsibilities. Asking them to become pioneers requires more than providing content. It requires a shift in identity and behaviour. Employees need to understand what is expected of them, why it matters, and how they can act within their local constraints.

The second tension concerns structure. The more decentralised a learning system becomes, the more participants need orientation. Feedback across levels repeatedly asks for clearer agendas, better expectation management, more concrete examples, more time for exercises, clearer next steps, and stronger guidance in group work. This does not contradict the value of peer learning. Rather, it clarifies the conditions under which peer learning works. Participants appreciate exchange and autonomy, but they also want a reliable frame.

This finding mirrors the literature. Participatory AI implementation is not the absence of structure. Ruess et al. emphasize that participation requires procedural models, infrastructural support, and phased involvement. Rueckert et al. similarly show that experimental spaces need deliberate design; simply exposing employees to AI is not sufficient (Ruess et al., 2024; Rueckert et al., 2021). The AI Pioneers case confirms this at enterprise scale. Peer learning requires structure precisely because it mobilizes heterogeneous participants across different contexts.

The third tension concerns engagement over time. Participation data shows a typical funnel from broad entry-level participation toward smaller advanced cohorts. This is not inherently problematic. Advanced levels naturally require more commitment, and not every participant needs to become a strategic AI project leader. However, drop-off becomes an issue if participants leave without transferring knowledge into their teams or without continuing to engage in the community. The important question is therefore not only how many employees attend modules, but how many become active carriers of AI capability.

The fourth tension concerns heterogeneity. Participants differ in AI experience, functional background, learning expectations, seniority, and access to implementation opportunities. A format that is accessible to beginners may feel too basic for advanced participants. A module that focuses on project management may be useful for some, but too distant from concrete AI tool use for others. Feedback from Level 2 and Level 3 shows this clearly: participants value methods and exchange, but often ask for more depth, more technical specificity, or clearer links to implementation.

This heterogeneity is unavoidable in large-scale transformation. The solution is not to eliminate it, but to design for it. The program's level structure is one response. Another is clearer expectation management: participants need to know whether a module focuses on orientation, experimentation, use-case development, governance, project implementation, or strategic reflection. Without this clarity, the same session may be perceived as valuable by some and too generic by others.

The fifth tension concerns governance and responsibility. AI transformation cannot be treated only as enablement. Employees also need to understand legal, ethical, and organisational guardrails. The Promptathon paper already showed that Deutsche Telekom integrated responsible AI and digital ethics into experiential learning formats (Kohn and Leffler-Krebs, 2025). The AI Pioneers program continues this logic.

This governance dimension strengthens the program, but it also complicates it. Employees want freedom to experiment, but they also need clarity on which tools are approved, which data may be used, which use cases require review, and how compliance requirements shape implementation. The practical challenge is to communicate governance not as a barrier to innovation, but as an enabling condition for responsible scaling.

7 Core insight: Peer learning needs community management

The central insight from the case is that peer learning does not scale autonomously. This may appear obvious, but it is often underestimated. Organisations sometimes assume that once a community exists, knowledge will naturally flow; once multipliers are trained, they will automatically multiply; once a learning journey is launched, transformation will continue by itself. The AI Pioneers case suggests otherwise.

Peer learning creates potential energy. It activates intrinsic motivation, distributed expertise, and cross-functional exchange. But this energy dissipates without orchestration. Community management is the organisational capability that keeps the learning system alive. It provides continuity, maintains engagement, curates content, connects people, clarifies expectations, supports follow-up, and translates decentralised learning into visible outcomes.

Community management is necessary for at least four reasons. First, it sustains motivation. Voluntary programs compete with daily work, deadlines, restructuring, managerial priorities, and personal workload. Even motivated participants may disengage if the community does not create rhythm, relevance, and recognition. Second, community

This paper was presented at the XXXVII ISPIM Innovation Conference, held in Granada, Spain on 07 June to 10 June 2026. ISBN 978-952-65771-2-8.

management provides orientation. Participants need to understand where they are in the learning journey, what is expected at each level, and how they can apply what they learn. Third, it supports knowledge transfer. Useful insights from one group need to be captured, shared, and made accessible to others. Fourth, it helps connect learning with implementation. Without follow-up, many ideas remain ideas. Community management can help identify owners, connect them to experts, and maintain momentum.

This finding refines the promise of peer learning. The answer to the research question is therefore not simply that peer learning enables scalable AI transformation. More precisely, peer learning enables scalable AI transformation when it is embedded in a structured community architecture that balances autonomy with orchestration.

This balance is important. Too much central control would undermine the local relevance and intrinsic motivation that make peer learning powerful. Too little structure would lead to fragmentation, passivity, and inconsistent outcomes. The managerial challenge is to create guided decentralisation: enough freedom for local exploration, but enough structure to make learning cumulative and actionable.

In innovation management terms, the AI Pioneers program can be understood as a capability-building infrastructure. It does not replace central AI strategy, governance, or technical platforms. Instead, it complements them by creating distributed human capability. It helps the organisation sense AI opportunities locally, mobilize employees around use cases, and connect learning with responsible implementation.

8 Discussion: Design principles for scaling AI transformation

The case provides several design principles for organisations seeking to scale AI transformation through peer learning. These principles are not a universal blueprint, but they are transferable as managerial guidelines.

The first principle is to decentralise capability building without abandoning central responsibility. AI opportunities are embedded in local work contexts. Therefore, capability building must happen close to teams, processes, and domain problems. However, decentralisation does not mean leaving employees alone. Central teams remain necessary for governance, tools, learning architecture, community management, and strategic alignment.

The second principle is to build on existing context competence. Organisations often search for AI experts externally or concentrate AI knowledge in technical units. The case suggests that a complementary approach is needed. Many employees already understand their domains deeply. AI transformation becomes more scalable when this domain knowledge is activated and combined with basic AI literacy. This principle is consistent with Pfeiffer's argument that employees' context competence is an underestimated foundation for AI implementation (Pfeiffer, 2020).

The third principle is to make learning experiential. AI cannot be understood fully through slides. Employees need to interact with tools, test prompts, compare outputs, evaluate limitations, and apply AI to real tasks. This is consistent with both the Promptathon experience and the broader literature on participatory and experiential AI learning (Kohn and Leffler-Krebs, 2025; Rueckert et al., 2021; Schoenmaker, 2021).

The fourth principle is to use peer networks as diffusion channels. In large organisations, knowledge does not scale only through formal communication. It scales through trust, proximity, relevance, and social proof. Employees are often more willing to try AI when they see colleagues in similar roles using it productively. Peer networks can therefore reduce uncertainty and make AI adoption more credible.

The fifth principle is to define and support multiplier roles. It is not enough to train individuals. Organisations need to clarify how these individuals should contribute after training. Are they expected to support colleagues, moderate workshops, identify use cases, act as governance contact points, or lead projects? The role must be explicit enough to guide action, but flexible enough to fit different contexts.

The sixth principle is to integrate responsible AI from the beginning. AI literacy is incomplete if it focuses only on productivity and tools. Employees also need to understand ethical, legal, and organisational boundaries. The Deutsche Telekom case shows that responsible AI can be embedded in practical learning formats rather than treated as a separate compliance module. This integration is especially important as regulatory expectations around AI literacy and responsible use increase.

The seventh principle is to invest in community management as a core transformation function. Community management should not be seen as administrative support. It is the mechanism that turns a collection of motivated individuals into a functioning learning network. It requires time, ownership, facilitation skills, communication routines, content curation, and management support.

Together, these principles suggest that AI transformation at scale is neither a pure training problem nor a pure technology rollout. It is an organisational learning challenge. The organisation must create structures through which employees can learn, experiment, share, and act responsibly.

9 Practical implications

For practitioners, the case offers several implications. First, organisations should avoid treating AI literacy as a one-off training campaign. Initial events such as Promptathons are valuable because they create awareness, confidence, and enthusiasm. However, sustained transformation requires follow-up structures that help employees transfer learning into their teams. The AI Pioneers program represents such a follow-up structure.

Second, organisations should design AI learning journeys with differentiated levels. Broad entry formats are necessary to lower barriers and create orientation. Advanced formats are needed for use-case development, implementation, and strategic reflection. A single format cannot serve all employees equally well. Level-based designs help manage heterogeneity, but only if expectations are clearly communicated.

Third, organisations should measure carefully and honestly. Internal satisfaction and usefulness ratings are valuable for steering a program, but they should not be overstated as proof of impact. Organisations should distinguish between participation, satisfaction, perceived usefulness, behaviour change, use-case implementation, productivity impact, and cultural transformation. The AI Pioneers data provides useful descriptive evidence, but future evaluation could be strengthened through follow-up surveys, use-case tracking, qualitative outcome stories, and indicators of team-level diffusion.

Fourth, organisations should recognize the workload of transformation. Peer learning may sound efficient because it distributes effort, but it still requires resources. Community management, facilitation, content development, governance alignment, and participant support all require time. If these tasks are added informally to already busy employees, the program risks depending too heavily on intrinsic motivation and a small number of highly committed individuals.

Fifth, leadership support remains important, but not sufficient. Leaders can legitimize the initiative, provide resources, and encourage participation. However, the everyday success of peer learning depends on local managers allowing employees time to participate, apply AI, and share knowledge. Without such local support, the program may be appreciated but not translated into practice.

Finally, organisations should treat peer learning as part of AI governance. Responsible AI use cannot be ensured only through policies. Employees need practical understanding of how governance applies to real work. Peer learning communities can help translate abstract rules into situated judgement, especially when connected to compliance, ethics, and tool guidance.

10 Conclusion

This paper examined how peer learning can enable scalable AI transformation in large organisations through the case of Deutsche Telekom's AI Pioneers program. The case shows that peer learning supports AI transformation by enabling localized knowledge creation, peer-driven diffusion, action-oriented learning, role-based multiplication, and community-based capability building.

The findings also show that peer learning is not self-sustaining. Decentralised learning creates relevance and engagement, but it also produces coordination challenges. Participants need structure, guidance, expectation management, and support for translating learning into action. The central contribution of the paper is therefore the insight that scalable peer learning requires active community management. AI transformation at scale depends on guided decentralisation: empowering employees locally while orchestrating the learning system centrally.

For innovation management practice, the case demonstrates that AI transformation should be understood as a social and organisational learning process. Tools matter, but they do not transform organisations by themselves. Transformation occurs when employees develop the confidence, competence, and responsibility to apply AI in their own contexts and when organisations create the structures that allow this learning to spread.

Future research could build on this case by comparing peer-learning approaches across organisations, developing stronger evaluation models for AI literacy initiatives, and examining how community-based learning translates into measurable business outcomes over time. For practitioners, the immediate lesson is clear: scaling AI transformation requires more than training more people. It requires building a community-enabled capability system that helps employees learn from each other, act responsibly, and turn AI potential into everyday organisational practice.

References and Notes

Antweiler, D., Albiez, D., Bures, D. et al. (2023) ‘Einsatz von KI-basierten Anwendungen durch Krankenhauspersonal: Aufgabenprofile und Qualifizierungsbedarfe’, Bundesgesundheitsblatt – Gesundheitsforschung – Gesundheitsschutz, 67, pp. 66–75. doi:10.1007/s00103-023-03817-x

Kohn, S. and Leffler-Krebs, S. (2025) ‘Building AI Literacy through Promptathons: The Deutsche Telekom Case’, XXXVI ISPIM Innovation Conference Proceedings, Bergen, Norway.

Long, D. and Magerko, B. (2020) ‘What is AI literacy? Competencies and design considerations’, Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1–16. doi:10.1145/3313831.3376727

Ng, W., Leung, J. K., Chu, S. and Qiao, M. S. (2021) ‘Conceptualizing AI literacy: An exploratory review’, Computers and Education: Artificial Intelligence, 2, 100041. doi:10.1016/j.caeai.2021.100041

Pfeiffer, S. (2020) ‘Kontext und KI: Zum Potenzial der Beschäftigten für Künstliche Intelligenz und Machine Learning’, HMD Praxis der Wirtschaftsinformatik, 57, pp. 465–479. doi:10.1365/s40702-020-00609-8

Rueckert, P., Papenberg, B., Sievers, T. and Tracht, K. (2021) ‘On-Site Learning Factories for Implementing Artificial Intelligence Systems in Production’, Proceedings of the Conference on Learning Factories (CLF). doi:10.2139/ssrn.3864445

Ruess, P., Staffa, A., Kreutz, A. et al. (2024) ‘Künstliche Intelligenz in betrieblichen Prozessen: Ein Vorgehensmodell zur partizipativen Gestaltung von KI-Anwendungen’, HMD Praxis der Wirtschaftsinformatik, 61, pp. 485–502. doi:10.1365/s40702-024-01049-4

Schoenmaker, F. A. M. (2021) ‘Teams an digitale Technologien heranführen’, BWK Energie, 73(3–4), pp. 24–25. doi:10.37544/1618-193X-2021-3-4

Declaration on the Use of AI Tools

In preparing this paper, AI-based tools were used as supporting instruments. Elicit supported the literature search, Assembly supported the transcription of interviews, and ChatGPT supported drafting, structuring, and language refinement. The authors reviewed, edited, and approved the final manuscript and remain fully responsible for its content, interpretation, and conclusions.