

Safeguarding Human Potential in Automated Work Systems: A Capability-Centred Framework for Design and Evaluation

LISA DIAMOND, AIT Austrian Institute of Technology, Austria

MICHAEL OPPERMANN, AIT Austrian Institute of Technology, Austria

Automation and AI increasingly reshape work systems beyond task execution, influencing how human capabilities are used, preserved, or transformed. Cognitive ergonomics traditionally evaluates automation through constructs such as workload, situation awareness, trust, and function allocation. While these perspectives explain important aspects of performance and interaction quality, they do not provide an integrated lens for evaluating how automation shapes human capabilities over time. This paper introduces the Human Potential in Automation (HPA) framework, a capability-centred evaluative framework that integrates perspectives from cognitive ergonomics, work design, and human-centred AI through four dimensions: Optimisation, Retention, Expansion, and Engagement. To operationalise the framework, we introduce a lightweight tool (“action cards”) to support participatory design reflection, and propose a concise 12-item companion scale to enable systematic evaluation of capability outcomes. We illustrate the framework through an exploratory multi-case study involving four industrial stakeholders in AI-supported quality control. The findings demonstrate the practical utility of the framework for identifying capability-related design considerations that extend beyond traditional performance-focused evaluation, contributing an integrative capability-centred perspective for the design and evaluation of automated work systems.

CCS Concepts: • **Human-centred computing** → **Human computer interaction (HCI)**; *HCI theory, concepts and models*; *Empirical studies in HCI*.

Additional Key Words and Phrases: human-automation interaction, cognitive ergonomics, capability sustainability, work design, human-centred AI

ACM Reference Format:

Lisa Diamond and Michael Oppermann. 2026. Safeguarding Human Potential in Automated Work Systems: A Capability-Centred Framework for Design and Evaluation. In *European Conference on Cognitive Ergonomics (ECCE 2026)*, September 15–18, 2026, Siena, Italy. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3822301.3822332>

1 Introduction

Automation has long been assessed through its effects on efficiency, workload reduction, and error mitigation. Foundational research has demonstrated that high levels of automation come with the possibility of introducing unintended side effects, including skill degradation, reduced situation awareness, limitation of skill transfer opportunities and diminished readiness for manual intervention [4, 14, 28]. Trust research has further highlighted how inappropriate reliance may lead to misuse or disuse of automated systems [23]. As AI-driven systems increasingly support perception, diagnosis, and decision-making, automation no longer merely reallocates manual control but reshapes the structure of professional expertise. This shift has been described as a transformation of how knowledge is distributed and applied

Authors’ Contact Information: Lisa Diamond, lisa.diamond@ait.ac.at, AIT Austrian Institute of Technology, Vienna, Austria; Michael Oppermann, michael.oppermann@ait.ac.at, AIT Austrian Institute of Technology, Vienna, Austria.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

Manuscript submitted to ACM

Manuscript submitted to ACM

within socio-technical systems, rather than a simple redistribution of tasks [21] and has led to increased attention to the role of humans in the loop, emphasising the need to maintain meaningful involvement in increasingly automated decision processes [12]. Contemporary research emphasises the tension between automation and augmentation, arguing that technology may either displace or amplify human capabilities depending on implementation choices [22, 30]. Empirical studies further highlight that effective human–AI collaboration depends on how systems support complementary human reasoning rather than passive reliance [5] and recent work further emphasises that human–AI collaboration involves the reconfiguration of roles, responsibilities, and knowledge boundaries within work systems, rather than a simple substitution or amplification of tasks [3]. Work design scholarship has similarly that digital systems can either undermine or reinforce autonomy, skill variety, and competence development [29]. From a psychological perspective, such changes are closely tied to professional identity, as shifts in task structure and autonomy affect how individuals understand and implement their roles within their work environments [32].

Despite these developments, evaluation practise in automation projects continues to prioritize short-term performance and adoption metrics. While explainability and transparency have increasingly been recognised as supporting human understanding and learning [38], evaluation remains largely centred on interaction quality and immediate performance with dominant constructs, including workload, situational awareness, trust, and function allocation providing explanations of specific aspects of interaction. Likewise, work design research focuses on autonomy and competence, while augmentation research examines complementary human–AI capabilities. These perspectives remain however fragmented and do not provide an integrated lens for evaluating whether automation preserves, develops, or undermines the human capabilities that organisations depend on over time for their successful survival and development. Consequently, designers must currently reason across multiple theoretical perspectives. To address this gap, we integrate these established perspectives into a coherent evaluative framework centred on long-term capability trajectories.

This paper introduces the Human Potential in Automation (HPA) framework, which conceptualises how automation shapes human potential along four interrelated dimensions: **Optimisation, Retention, Expansion, and Engagement** (the *HPA framework*). Rather than framing automation as inherently beneficial or harmful, the framework considers it as a design-dependent force that can enhance efficiency while either supporting or undermining human capabilities and engagement, depending on how work is structured and how human–automation interaction is shaped. To operationalise this framework, we introduce a set of action cards as a lightweight applied resource to help linking the dimensions to observable risk patterns and design interventions, and a proposed 12-item HPA scale to support the evaluation of capability-related outcomes of the human-automation cooperation. We illustrate the application of the framework through an exploratory multi-case study with four industrial stakeholders in the context of AI-driven quality control. In summary, we contribute:

- A **theoretical framework** (*HPA framework*) for mapping how automation affects human capabilities along four interrelated dimensions.
- An **applied tool** (*HPA action cards*) to support capability-centred reflection in automation design and review processes of already implemented solutions.
- A **proposed scale for validation** (12-item *HPA scale*) to complement established workload, trust, and engagement measures and evaluate both concepts and implemented workflows.
- An **exploratory multi-case study** illustrating how the HPA framework and its instruments can support requirements elicitation, design reflection, and evaluation of human–automation systems in industrial practise.

The remainder of this paper is organized as follows. Section 2 reviews related work across cognitive ergonomics, automation research, and work design. Section 3 presents the HPA framework and its four dimensions. Section 4 describes the action cards and Section 5 introduces the proposed HPA scale. Section 6 reports the multi-case study. Section 7 discusses implications and limitations.

2 Background and Related Work

2.1 Cognitive Ergonomics and Human–Automation Interaction

Research in human–automation interaction has provided explanations of how task allocation influences cognitive load, error likelihood, and takeover performance [13, 28]. These models primarily address immediate interaction dynamics but do not consider long-term trajectories of human expertise within automated systems. While related work has identified phenomena such as the out-of-the-loop performance problem and reduced operator engagement under high levels of automation [14], these effects are typically treated as interaction-level issues rather than as part of broader capability dynamics. This limitation is also reflected in cognitive systems engineering literature, which emphasises adaptive capacity and the ability of human–system collaborations to respond to changing demands over time, but does not provide explicit means to assess how such dynamics shape human capabilities [41].

2.2 Augmentation, Deskilling, and Capability Trajectories

Augmentation perspectives, originating in early visions of intelligence amplification [15], propose that computing technologies can expand human cognitive capability rather than substitute it. This tension has been described as the automation–augmentation paradox: the same system can reduce or enhance human capability depending on how it is implemented [30]. Human-centred AI also frames this as an ethical imperative: systems should be designed to amplify rather than replace human judgement [35] and evidence suggests that human–AI performance improves when systems promote active reasoning rather than passive use [5]. Yet augmentation remains largely normative and availability of systematic instruments to assess whether augmentation actually occurs in practise is limited.

2.3 Work Design and Professional Agency

Work design theories emphasise that autonomy and skill utilization are central to sustained motivation and well-being [18]. Contemporary extensions highlight the particular importance of these factors under conditions of algorithmic coordination, where system design could inadvertently narrow task scope or reduce decision spaces [29]. Self-determination theory identifies autonomy and competence as fundamental psychological needs; automation that undermines either of them risks eroding both motivation and professional identity [9]. At the same time, work design research shows that employees actively shape their work through practises such as job crafting, suggesting that capability development depends not only on system constraints but also on opportunities for active engagement throughout the work process [42].

2.4 Positioning

The HPA framework emerged through an iterative synthesis of literature across cognitive ergonomics, human-centred AI, and work design, combined with recurring capability-related concerns that were identified through stakeholder engagement in multiple industrial automation projects from different contexts. The resulting four HPA dimensions organize these complementary perspectives into a coherent capability-centred framework for evaluating automation.

Table 1. Capability trajectories in automated work systems: the four HPA dimensions.

Dimension	Core focus	Capability risk <i>if misaligned</i>	Capability potential <i>if well designed</i>	Example design levers
Optimisation	Effective use of human strengths	Underutilization of human potential; monitoring-only roles	Focus on high-value judgement and complex problem-solving	<i>workload calibration; complementary function allocation</i>
Retention	Preservation of expertise	Skill decay; erosion of tacit knowledge; loss of practise	Sustained expertise; continued knowledge sharing	<i>practise intervals; trust calibration; override mechanisms; knowledge capture and transfer</i>
Expansion	Growth of capability	Skill stagnation; passive tool reliance	Learning; new insights; broadened task scope	<i>explainability design; feedback loops; data access</i>
Engagement	Professional agency and motivation	Reduced autonomy; loss of professional identity	Sustained autonomy; active judgement; positive collaboration experience	<i>task variation; purpose transparency; human control scope</i>

The HPA framework shifts the analytical focus from isolated interaction aspects toward their integrated implications for human capability. It examines how automation design decisions collectively shape the use, preservation, development, and application of human capability over time, providing a complementary evaluative perspective on whether automation sustains, transforms, or diminishes the capabilities within the work force that organisations depend on in the long term.

The HPA action cards and proposed scale were subsequently developed to operationalise the framework for participatory reflection to identify risks and opportunities and structured evaluation (following scale validation).

3 The Human Potential in Automation Framework

Human potential is defined here as the collection of skills, experience-based knowledge, professional judgement, and motivational engagement that individuals bring to and develop through work. Automation reshapes this potential not only by reallocating tasks, but by influencing how capabilities are practised, valued, expanded, or marginalized over time. In AI-enabled work systems, this increasingly involves a shift in how expertise is organized and applied within the system [3, 21].

Research in human–automation interaction has long demonstrated that design decisions influence operator readiness, situation awareness, and control performance [13, 28]. However, these models primarily focus on performance states. The HPA framework shifts attention toward *capability trajectories*, conceptualising automation as generating degradation risks or development opportunities depending on alignment between system design and human strengths. Within this perspective, optimisation and engagement frame how capabilities are utilized and experienced, while retention and expansion capture how capabilities are maintained and developed over time. This perspective aligns with broader evidence that technological systems reshape how human capabilities are deployed, valued, and sustained over time, rather than simply reducing or enhancing them [1].

Drawing on the preceding synthesis of cognitive ergonomics, work design, and human-centred AI literature, together with recurring capability-related concerns identified through initial stakeholder engagement, four complementary dimensions emerged that frame and structure these trajectories: **Optimisation, Retention, Expansion, and Engagement**. Together, they represent the primary capability concerns consistently identified across theory and practise and provide a structured lens for capability-centred conceptualisation and evaluation, while allowing for further refinement as the framework evolves. Table 1 provides an overview of each dimension’s core focus, associated capability risks, potential opportunities when system design is well-aligned, and example design levers to inform design responses.

3.1 Human Potential Optimisation

Optimisation concerns whether automation enables effective alignment between human strengths and task demands. Function allocation theory emphasises that performance improves when tasks are distributed according to the complementary abilities of humans and machines [28, 34]. When automation is introduced primarily for efficiency without regard to human strengths, workers risk be reduced to supervisory or monitoring roles, leading to underutilisation of expertise. From a work design perspective, meaningful skill utilisation is central to sustained motivation [18], revealing a close coupling between the optimisation and engagement dimensions. Recent work on hybrid intelligence further emphasises that effective human–AI systems depend on aligning system capabilities with human strengths such as contextual reasoning and sensemaking [10]. Optimisation therefore captures whether automation creates a trajectory toward continued, optimal capability utilisation.

3.2 Human Potential Retention

Retention addresses whether existing expertise remains practised and relevant despite automation. Early research demonstrates how removing operators from routine engagement, leaving them responsible only for rare, high-consequence interventions, increasing the risk of skill degradation and consequent errors [4]. The out-of-the-loop problem further illustrates how reduced active involvement weakens readiness and procedural competence [14]. In contemporary digital systems, algorithmic centralization of decision logic can similarly diminish distributed professional expertise [22]. At the same time, studies on trust calibration suggest that reduced interaction and overreliance on automated outputs can further weaken active engagement and skill maintenance over time [24, 39]. Retention therefore concerns the continuity of practised competence and knowledge transmission. A positive trajectory is achieved when automation preserves opportunities for active judgement and sustained expertise use and transfer.

3.3 Human Potential Expansion

Expansion captures whether automation broadens cognitive scope and enables capability growth. Early visions of computing as a means of augmenting human intellect emphasised the potential for perception expansion and analytical capacity rather than merely automating routine tasks [15]. Recent discussions of augmentation versus automation argue that systems can deepen expertise by surfacing patterns, counterfactuals, or alternative interpretations that support reflective judgement [30]. Previous work further shows that collaboration quality depends on whether systems invite active interpretation and judgement, instead of encouraging uncritical acceptance of outputs. [5]. At the same time, explainability and transparency have been identified as key conditions for enabling understanding and learning in human–AI interaction [38]. When systems operate as black boxes or reduce user involvement to passive acceptance, capability can stagnate. By contrast, interaction design elements such as reliability indicators can support more calibrated use, active engagement with system outputs, and learning from the loop, influencing how users interpret and act on automated recommendations [17]. Expansion thus reflects whether automation creates trajectories of capability growth such as learning, insight generation, and access to new task domains or, if misaligned, leads to skill stagnation and limited development.

3.4 Human Potential Engagement

Engagement concerns whether interaction with automation supports professional agency and motivational sustainability. Self-determination theory identifies autonomy and competence as fundamental psychological needs [9] and

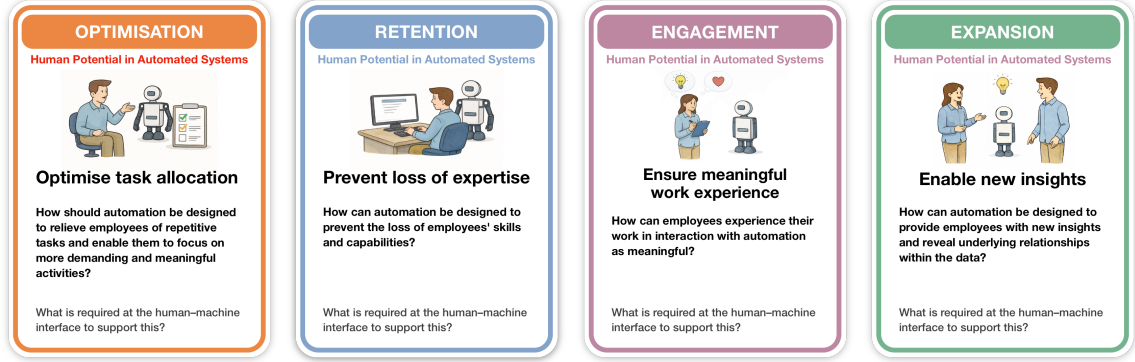


Fig. 1. Exemplary action cards for the four HPA dimensions (optimisation, retention, expansion, engagement). Each card operationalises a dimension through a design-oriented principle and guiding question to support reflection in participatory design and evaluation settings.

Work design research shows that digital systems can either undermine or reinforce these needs depending on implementation [29]. From a broader perspective, shifts in autonomy and task structure are closely linked to professional identity, influencing how individuals understand and perform their roles in AI-mediated work environments and experience meaningfulness with a critical impact on sustained engagement and well-being [32]. Designs that centralize decisions, reduce opportunities for professional judgement application, or obscure responsibility boundaries risk eroding professional identity and intrinsic motivation over time. Designs that preserve decision autonomy and make human judgement visible, in contrast, enable trajectories of sustained agency and engagement.

4 Action Cards

To facilitate application of the framework in practice, we introduce a lightweight action card format linking each dimension to observable opportunity and risk patterns and targeted design interventions. Card-based approaches have demonstrated broad utility as structured design tools to support reflection, creativity, and engagement in learning and design contexts (e.g., [8, 11, 25, 43]). The proposed HPA action cards serve as guiding prompts in participatory design or review workshops, helping teams identify early indicators of capability degradation, formulate concrete adjustments in task allocation, interface design, or governance structures, and reflect on capability sustainability throughout the automation lifecycle.

Three action cards were designed per HPA dimension, forming a set of 12 cards. Each action card addresses one HPA dimension and combines a dimension-specific focus statement with guiding questions that prompt reflection on potential risks and design responses. Figure 1 shows one exemplary card for each of the four HPA dimension. Each card includes the HPA dimension name, an illustration, a concise design-oriented principle (e.g., ensuring meaningful work experience, optimising task allocation, preventing the loss of expertise, enabling new insights), and a reflective question aiming to help teams in considering these aspects.

The full set of action cards in high-resolution (PDF) is available as supplemental material at: <https://osf.io/uejfw/>

Table 2. The HPA scale: 12 items across four dimensions.

Dimension	Nr	Item
Human Potential Optimisation	1	The process automation enables employees to carry out their work more efficiently.
	2	The process automation helps employees achieve high-quality job performance.
	3	Automation relieves employees of routine tasks and allows them to make optimal use of their skills.
Human Potential Retention	4	Employees' skills remain important despite the process automation.
	5	The process automation helps maintain the know-how and skills of employees.
	6	Employees' expert knowledge continues to be shared and used within the automated process.
Human Potential Expansion	7	The process automation enables employees to gain additional insights that support their expertise.
	8	Employees can learn new skills or techniques through interacting with the automated process.
	9	The process automation facilitates employees carrying out new tasks.
Human Potential Engagement	10	The process automation supports rather than replaces professional judgement and expertise.
	11	Employees retain autonomy and decision-making freedom despite the process automation.
	12	Employees experience their interaction with the automated process as positive.

5 The HPA Scale

To support empirical assessment of capability-related outcomes across the four HPA dimensions, we propose a complementary 12-item scale structured around the four HPA dimensions.

Three items operationalise each dimension's core theoretical construct into evaluative statements at the level of "process automation," enabling domain-independent application across industrial, service, and knowledge work contexts (with the option to adjust the phrasing in order to refer to a specific process). They are designed for process-level assessment (e.g., by managers) but can be adapted for individual user perspectives. The scale is intended for diagnostic use during automation pilots, comparative evaluation of design alternatives, and longitudinal monitoring of perceived capability sustainability. The proposed HPA scale complements established measures of workload, trust, and work engagement (e.g. NASA-TLX [19], TOAST [40], NSa-WS/NSu-WS [37]), by providing a focused evaluation layer addressing how automation affects capability use, preservation, development, and agency.

Table 2 presents the full item set. Items use a Likert-type response format (1 = strongly disagree, 5 = strongly agree).

6 Multi-Case Exploratory Study: AI-Assisted Industrial Quality Control

We conducted a multi-case exploratory study with four industrial stakeholders, each operating in a sector characterized by increasing automation in quality control: aerospace components, high-performance robotic manufacturing, specialty paints, and coatings. The four organisations were selected to capture a range of settings in which automation affects capabilities. They span different industrial areas, differ in automation maturity and in their primary orientation, with one organisation focused on research and development and the others on production. Across the four organisations, six representatives from management, implementation, and research & development, who actively influence implementation of automation processes, served as our main contacts.

All four organisations shared a common set of challenges: deploying AI-driven visual defect detection while avoiding the loss of human expertise; enabling workers to recognize when and why the automated assessment fails; developing trust calibration capabilities; and fostering structured documentation practises that enable insights and support training. The study was conducted as a collaborative requirements analysis and system evaluation exercise, with HPA constructs used to structure inquiry and reflection across phases.

6.1 Procedure

The study followed an adapted design study methodology [31], oriented toward requirements elicitation and capability-centred design reflection. We began with a joint workshop involving all four stakeholder organisations, where we introduced and discussed the HPA dimensions in the context of automation workflows and visual inspection needs for AI-driven industrial quality control. We then conducted individual follow-up sessions with each stakeholder to explore context-specific challenges and interaction-oriented opportunities facilitated through the lens of the HPA framework dimensions, using the proposed action cards and associated interview questions derived through the HPA framework. This was followed by feedback gathering on a conceptual model of the overall process and initial interface concepts based on the insight from the joint workshop. Following the development of a demonstration application (Sec. 6.2) and internal pilot testing, we held individual 1.5-hour evaluation sessions with each stakeholder. Using a chauffeured demo format, an established practise in visualization research [6, 26], we guided participants through the system and its key functionalities using the developed technology probe (see below) and collected qualitative feedback along the HPA dimensions and other user experience and impact oriented aspects. At the end of each session, we shared a questionnaire — including the HPA scale — which participants completed independently to capture perceived potential capability impacts of the AI quality control workflow. Importantly, this study was not intended as a usability evaluation of an interface, but to assess participants’ broader reflections on how automated quality control may affect capabilities and roles in their organisations and how interface design choices may or may not succeed in making use of opportunities and mitigating risks arising regarding human potential preservation and utilization in process automation.

6.2 Technology Probe

To enable consistent evaluation across all four cases, and to avoid dependency on proprietary production data and workflows, we developed a demonstration application using an open-source solar cell defect dataset (PVEL-AD [36]). It functioned as a technology probe [20]: a purposefully simplified inspection tool to place stakeholders in a realistic human–automation context that reflects their task of AI-driven visual defect detection. Importantly, we used real defect data as it has been shown to support engagement and meaningful reflection in design study contexts [27].

The application (see Fig. 2) provides both an overview for aggregate analysis (expansion, optimisation) and a detail view (optimisation) for the inspection of individual samples. The overview (Fig. 2a) supports the analysis of patterns across samples, shows distributions (e.g., number of human corrections), and allows users to interpret results within a broader production context (expansion).

The detail view (Fig. 2b) presents the annotated photo alongside model outputs and supports users in dual roles as inspectors and knowledge contributors (engagement, retention). Correction and curation features allow users to apply tags and notes, override classifications, and select representative samples for inclusion in a shared knowledge base (optimisation, retention, engagement). These interactions make individual assessments persistent and reusable.

In addition to the predefined views, natural language querying [33] provides an ad hoc analytical layer for context-specific questions (optimisation, extension). Queries may include retrieving defect counts within a time range or viewing samples in a specific defect category. The goal is to lower the barrier for workers in operational roles whose responsibilities increasingly extend to data interpretation and process monitoring.



Fig. 2. The demo application, designed to represent shared aspects of AI-assisted quality control. (a) The overview page provides aggregated visualizations, including temporal defect patterns, a spatial heatmap of defect locations, and a sample gallery for collection-level review. (b) The detail page presents an annotated image, a defect score, and features supporting knowledge transfer: tagging, notes, manual classification override, and knowledge base assignment. Natural language querying of the collected data is available across the application.

6.3 Analysis

We captured and organized the collected qualitative and quantitative data systematically across all phases. The 2-hour joint workshop used a shared digital whiteboard (Miro) to elicit and cluster stakeholders' observations. The individual follow-up and evaluation sessions were audio- and screen-recorded, complementing the questionnaire responses. As the sample of four stakeholder organisations was too small to support meaningful quantitative analysis, we used the questionnaire responses as additional reflective input, supporting a deeper understanding of the qualitative insights. Two members of the research team reviewed the responses, and the derived themes under each HPA dimension were then discussed jointly by the authors. We followed an interpretive approach by organizing observations under the four HPA dimensions and comparing them across the four cases (S1–S4) to identify recurring and divergent concerns.

6.4 Findings

We now report our synthesized, qualitative insights from workshop discussions and individual sessions across all four cases (S1–S4). Stakeholders engaged meaningfully with the framework, and each HPA dimension yielded distinct and substantive findings, regardless of the organisational context. The dimensions surfaced non-overlapping, complementary concerns and suggest that the framework provides adequate coverage to conceptualise human potential in automation. Across the four cases, the HPA framework enabled a structured analysis of capability-related implications that extended beyond traditional performance considerations. Rather than discussing automation solely in terms of efficiency or usability, stakeholders consistently reflected on how system design would influence the use, preservation, development, and sustained exercise of human capabilities. Although organisations placed different emphasis on individual capability dimensions, discussions consistently highlighted the need to balance multiple capability considerations when designing and implementing automated processes.

Optimisation: All four organisations viewed automation as an opportunity to improve efficiency and precision in quality control, though the specific gains they anticipated differed. Common themes included reducing resource losses by catching defects before further processing (S3), earlier identification of lacking product homogeneity within the production cycle (S2), reduction of subjective variability in quality judgements through standardized measurement (S1), and freeing workers from routine error-searching to focus on higher-level problem-solving (S4). A distinction was

observed between research-oriented and production-facing contexts. Stakeholders from R&D-focused organisations (S1, S2) emphasised optimisations regarding quality and thoroughness, while those in production settings (S3, S4) additionally focused on speed of defect identification. In both contexts, automation was expected to augment rather than replace expert judgement.

Retention: Preserving human knowledge was articulated as a concern across all organisations, though the perceived urgency varied. A recurring point was that expertise cannot be fully delegated to automated systems: workers need to retain sufficient understanding of the process to verify model outputs, recognize failures, and intervene when necessary. Attention decay was identified as a concrete downstream risk; as reliance on automated outputs increases, workers' readiness to detect anomalies outside the model's scope may erode over time (S1, S3).

Expansion: Automation was perceived as an opportunity to extend analytical capabilities and support knowledge documentation. Rather than replacing expert judgement, shifts in where that judgement is applied may occur, e.g. toward monitoring and evaluating system performance. Stakeholders pointed to the potential for systematic data collection to enable analyses not feasible in manual inspection (S1, S3, S4), the use of historical defect records for training and longitudinal benchmarking (S1, S3), and feedback mechanisms that make error patterns more visible (S3). The technology probe also prompted reflection on onboarding. Several stakeholders suggested that this type of interactive visualization alongside automation could enable less experienced employees to work on analytical tasks and thereby widen participation in quality processes.

Engagement: All stakeholders expressed worker acceptance and sustained motivation as a central design challenge. Automated quality control systems need to be comprehensible, meaningful, and appropriately demanding to support long-term workforce acceptance. A recurring theme was the need to make the purpose of automation-related responsibilities, particularly documentation and data curation, transparent and meaningful rather than perceived as administrative overhead (S4). At the same time, the potential of automation to reduce physically demanding and monotonous inspection tasks was seen as a source of motivational benefit (S1, S3). The temporal structure of automated workflows was identified as a key factor for sustained engagement. Deliberate task variation and planned interruptions during automated process steps were discussed to prevent passive monitoring and maintain active involvement over time (S1, S3). Interface design was considered a support mechanism to help workers maintain orientation within complex, multi-step inspection processes. S4 also noted exploratory interest in gamified interaction elements, as a way to create meaningful feedback loops within routine quality control work.

Taken together, the findings demonstrate that the four HPA dimensions captured distinct but interrelated capability implications of automation across all four cases, with automation consistently conceptualised as supporting human work rather than replacing it. Optimisation was predominantly associated with improving quality, consistency, and enabling experts to focus on higher-value problem-solving; retention was consistently framed as preserving the knowledge required to critically supervise automation, validate automated outputs, and intervene when necessary; expansion was primarily viewed as creating opportunities for learning, knowledge documentation, and broader participation in analytical tasks; engagement highlighted that sustained acceptance depended on preserving meaningful responsibility, task variety, and a clear purpose for automation-supported work. The findings also illustrate that these dimensions were frequently discussed in combination, suggesting that practitioners naturally reason about automation as balancing multiple capability considerations rather than optimising a single objective. This indicates that the HPA framework provides a useful analytical structure for examining how automation design simultaneously shapes different aspects of human capability.

7 Discussion

7.1 From Interaction Quality to Capability Trajectories

While existing guidelines for human–AI interaction emphasise usability, feedback, and appropriate reliance, they do not explicitly address how design decisions shape human capabilities over time [2] and related empirical work shows that the effectiveness of human–AI systems depends not only on system accuracy, but on how interaction design supports complementary human reasoning and judgement, rather than encouraging passive reliance [5]. This aligns with broader perspectives on human–AI collaboration that emphasise the complementary strengths of humans and AI as a central design challenge. [10] and connects to the emerging regulatory context for automated work systems: the EU AI Act [16] establishes human oversight as a requirement for high-risk AI systems, and such oversight presupposes that workers retain the expertise and agency to understand and intervene in automated decisions.

Against this background, the HPA framework contributes to cognitive ergonomics by introducing a capability-centred evaluation lens that complements established constructs. Traditional measures such as workload, trust, and situational awareness remain essential for assessing interaction quality in automated systems but are largely state-based and do not directly capture whether automation sustains, transforms, or redistributes professional capability over time. By conceptualising automation as shaping capability trajectories, the framework brings together perspectives on deskilling, augmentation, professional autonomy, and skill utilisation within a common evaluative lens. It thus supports a shift from short-term performance optimisation toward long-term capability sustainability as a design and evaluation aspect in automated work systems. This reframing is particularly relevant as AI systems increasingly operate not as discrete tools but as pervasive organisational infrastructure shaping how expertise is practised and transmitted.

7.2 Implications for Design and Evaluation Practise

The multi-case exploratory study presented demonstrates that HPA constructs are meaningful and accessible to industrial practitioners. The action cards provide a productive means of surfacing implicit assumptions about capability effects during requirements analysis, enabling systematic reflection without requiring specialized expertise in cognitive ergonomics. They therefore function as a lightweight translation of cognitive ergonomics principles into design practise, making considerations such as skill use, situational understanding, and decision authority explicitly discussable in early development stages and translating insights into actionable guidance. In line with prior work, such structured reflection can help make otherwise implicit assumptions about human–AI interaction visible and actionable in design processes [2]. The technology probe and HPA scale complemented this qualitative process by enabling capability-oriented evaluation of concrete interface concepts. Although the limited number of participating stakeholders did not permit systematic psychometric validation of the scale, its application alongside the evaluation demonstrated its potential as a structured complement to qualitative assessment, making capability-related effects more explicit and comparable across design alternatives. Together, the case studies illustrate the feasibility of applying the HPA framework across successive stages of a capability-centred design process, from early requirements elicitation through interface evaluation. More broadly, this also responds to recent calls to move beyond performance-centric evaluation toward metrics that capture human–AI complementarity and long-term human outcomes [5].

The HPA framework is intended to complement existing cognitive ergonomics practises rather than replace them. Established design and evaluation methods—including workload assessment, trust measures, and usability evaluations—remain essential for understanding different aspects of human–AI interaction. Rather than replacing established approaches, the framework introduces a capability-centred perspective that can be applied alongside them throughout

the automation lifecycle, providing a structured lens to consider how design decisions influence the development, maintenance, and use of human capabilities, while the accompanying tools support both participatory reflection and structured evaluation. Applied iteratively throughout the process, these tools support organisations in monitoring capability-related outcomes over time and identifying emerging design issues requiring further refinement.

Beyond supporting iterative design and evaluation, the framework foregrounds capability sustainability as a legitimate design criterion, on par with efficiency and safety. This extends established cognitive ergonomics concerns—such as workload, situation awareness, and control—toward a more explicit consideration of how human capabilities are maintained and developed over time. In this sense, the framework complements existing work design perspectives by making the capability implications of system design decisions more directly assessable [29].

7.3 Scope, Limitations, and Future Research Directions

The HPA framework provides an initial conceptual basis for assessing human capabilities in automated work systems.

The multi-case exploratory study used a demonstration application as a technology probe and therefore findings reflect stakeholder perceptions and anticipated capability effects rather than longitudinal measurements of actual capability change. An important direction for future work is ongoing engagement with stakeholders across the automation lifecycle to examine how the framework supports reflection and evaluation over time and what additional insights into capability trajectories emerge in practise. Such longitudinal perspectives are particularly important given that human–AI collaboration and expertise development unfold through sustained interaction and organisational embedding rather than isolated episodes of use [3].

The proposed HPA scale has not yet been subjected to formal psychometric validation. Therefore, internal consistency, construct validity, and sensitivity to actual capability changes remain to be established [7] and a dedicated validation study needs to examine the scale’s psychometric properties and clarify its relationship to established performance measures.

Further, the action cards have been used as a reflective prompt during requirements analysis workshops but have not been evaluated as a standalone design intervention. Examining how well they enable non-expert teams to identify capability risks independently of the researchers remains an important research direction. In this regard, future work could also examine how specific design features—such as transparency or reliability indicators—support more calibrated and capability-supportive interaction with automated systems [17].

Lastly, the HPA framework should be understood as an initial conceptual synthesis rather than a definitive taxonomy of human capability in automated work. The four dimensions emerged through iterative integration of literature across cognitive ergonomics, work design, and human-centred AI, informed by recurring capability-related concerns identified through stakeholder engagement across multiple industrial automation projects. The four industrial cases share a common focus on AI-assisted visual inspection, providing preliminary evidence of the framework’s utility while leaving the generalisability of specific findings to other automation contexts as an open question. Future research should therefore examine both the conceptual scope and distinctiveness of these dimensions. In particular, broader dimensions such as engagement may benefit from further differentiation into more specific constructs, while studies in additional domains may reveal further capability dimensions or relationships that extend the current framework. Because workshops explicitly introduced the HPA dimensions before dedicated reflection activities, future work should investigate whether practitioners independently identify similar capability concerns when applying the framework without researcher facilitation. Applying the framework in other work settings could support examination

and refinement of its dimensions and help to identify additional design concerns not encountered in manufacturing contexts.

8 Conclusion

Automation should be evaluated in terms of how human capabilities evolve within it. The HPA framework, accompanying action cards, and proposed HPA scale provide an integrative, capability-centred approach for assessing and shaping capability trajectories in automated work systems. In doing so, they extend cognitive ergonomics toward a capability-centred perspective concerned with the long-term development and sustainability of human potential amid the growing reliance on AI and automation in work systems.

Engagement with industrial stakeholders across four sectors provides preliminary evidence that the framework supports structured reflection on capability-related implications of automation, helping identify capability risks and design opportunities that lie outside the scope of performance-focused evaluation. More broadly, the work contributes a vocabulary and evaluative lens for a central challenge in AI-augmented work: understanding and shaping how automation systems reconfigure, redistribute, and transform human expertise over time. Rather than framing this solely in terms of preservation or loss, the HPA perspective highlights the need to make capability trajectories visible and designable, ensuring that evolving human–AI systems remain aligned with the development, use, and sustainability of human capabilities.

In this sense, the framework shifts attention from evaluating systems as isolated technical artifacts toward understanding them as embedded elements of evolving work practises that shape how expertise is enacted, shared, and maintained. It highlights that decisions about automation are simultaneously decisions about the future configuration of human capability within organisations.

By making these dynamics explicit, the HPA framework supports more deliberate and reflective design processes during which capability implications are considered alongside efficiency and performance from the outset. This opens up new possibilities for designing automation not only as a means of optimisation, but as a driver of sustained learning, capability development, and resilient human–AI collaboration over time.

Acknowledgments

This work is part of the ZERO3 project (FO999896399), funded by the Austrian Research Promotion Agency (FFG), and was supported by the III.Lab (Intelligent Interfaces Innovation), funded by AIT Austrian Institute of Technology, University of Salzburg, and Land Salzburg (GA No. 20102/F2300416-KZP). The authors thank ADLER, TIGER Coatings, FACC, STIWA Automation, and the other project partners for their valuable contributions. Generative AI was used for grammar and language polishing and to support the creation of the action-card illustrations, with all outputs reviewed and verified by the authors.

References

- [1] Daron Acemoglu and Pascual Restrepo. 2018. Artificial intelligence, automation, and work. In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, 197–236. doi:10.3386/w24196
- [2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. 1–13. doi:10.1145/3290605.3300233
- [3] Callen Anthony, Beth A Bechky, and Anne-Laure Fayard. 2023. “Collaborating” with AI: Taking a system view to explore the future of work. *Organization Science* 34, 5 (2023), 1672–1694. doi:10.1287/orsc.2022.1651
- [4] Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19, 6 (1983), 775–779. doi:10.1016/0005-1098(83)90046-8

- [5] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the whole exceed its parts? The effect of AI explanations on complementary team performance. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. 1–16. doi:10.1145/3411764.3445717
- [6] Jürgen Bernard, Mara Solen, Helen Novak Lauscher, Kurtis Stewart, Kendall Ho, and Tamara Munzner. 2025. VIVA: Virtual Healthcare Interactions Using Visual Analytics, With Controllability Through Configuration. *IEEE Trans. on Visualization and Computer Graphics (TVCG)* (2025). doi:10.1109/TVCG.2025.3599458
- [7] Godfred O Boateng, Torsten B Neilands, Edward A Frongillo, Hugo R Melgar-Quinonez, and Sera L Young. 2018. Best practices for developing and validating scales for health, social, and behavioral research: a primer. *Frontiers in Public Health* 6 (2018), 149. doi:10.3389/fpubh.2018.00149
- [8] Liuqing Chen, Wengteng Cheang, Zhaojun Jiang, Yuan Xu, Zebin Cai, Lingyun Sun, Peter Childs, Ji Han, Preben Hansen, and Haoyu Zuo. 2025. I-Card: A Generative AI-Supported Intelligent Design Method Card Deck. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. 1–22. doi:10.1145/3706598.3713934
- [9] Edward L. Deci and Richard M. Ryan. 2000. The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry* 11, 4 (2000), 227–268. doi:10.1207/S15327965PLI1104_01
- [10] Dominik Dellermann, Philipp Ebel, Matthias Söllner, and Jan Marco Leimeister. 2019. Hybrid intelligence. *Business & Information Systems Engineering* 61, 5 (2019), 637–643. doi:10.1007/s12599-019-00595-2
- [11] Salma Elsayed-Ali, Sara E Berger, Vagner Figueredo De Santana, and Juana Catalina Becerra Sandoval. 2023. Responsible & Inclusive Cards: An online card tool to promote critical reflection in technology industry work practices. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. 1–14. doi:10.1145/3544548.3580771
- [12] Christos Emmanouilidis, Petros Pistofidis, Luka Bertonec, Vasilis Katsouros, Apostolos Fournaris, Christos Koulamas, and Cristobal Ruiz-Carcel. 2019. Enabling the human in the loop: Linked data and knowledge in industrial cyber-physical systems. *Annual Reviews in Control* 47 (2019), 249–265. doi:10.1016/j.arcontrol.2019.03.004
- [13] Mica R. Endsley. 2017. From here to autonomy: Lessons learned from human–automation research. *Human Factors* 59, 1 (2017), 5–27. doi:10.1177/0018720816681350
- [14] Mica R. Endsley and Evelyn O. Kiris. 1995. The out-of-the-loop performance problem and level of control in automation. *Human Factors* 37, 2 (1995), 381–394. doi:10.1518/001872095779064555
- [15] Douglas C. Engelbart. 1962. *Augmenting Human Intellect: A Conceptual Framework*. Technical Report. Stanford Research Institute.
- [16] European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series. Art. 14, human oversight.
- [17] Peter Fröhlich, Alexander G Mirnig, Damiano Falcioni, Johann Schrammel, Lisa Diamond, Isabel Fischer, and Manfred Tscheligi. 2022. Effects of reliability indicators on usage, acceptance and preference of predictive process management decision support systems. *Quality and User Experience* 7, 1 (2022), 6. doi:10.1007/s41233-022-00053-0
- [18] J. Richard Hackman and Greg R. Oldham. 1976. Motivation through the design of work: Test of a theory. *Organizational Behavior and Human Performance* 16, 2 (1976), 250–279. doi:10.1016/0030-5073(76)90016-7
- [19] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in Psychology*. Vol. 52. 139–183. doi:10.1016/S0166-4115(08)62386-9
- [20] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. 2003. Technology probes: inspiring design for and with families. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. 17–24. doi:10.1145/642611.642616
- [21] Mohammad Hossein Jarrahi. 2018. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons* 61, 4 (2018), 577–586. doi:10.1016/j.bushor.2018.03.007
- [22] Katherine C. Kellogg, Melissa A. Valentine, and Angèle Christin. 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14, 1 (2020), 366–410. doi:10.5465/annals.2018.0174
- [23] John D. Lee and Katrina A. See. 2004. Trust in automation: Designing for appropriate reliance. *Human Factors* 46, 1 (2004), 50–80. doi:10.1518/hfes.46.1.50.30392
- [24] Stephanie M Merritt, Deborah Lee, Jennifer L Unnerstall, and Kelli Huber. 2015. Are well-calibrated users effective users? Associations between calibration of trust and performance on an automation-aided task. *Human Factors* 57, 1 (2015), 34–47. doi:10.1177/0018720814561675
- [25] Novia Nurain, Chia-Fang Chung, Clara Caldeira, and Kay Connelly. 2024. Designing a Card-Based Design Tool to Bridge Academic Research & Design Practice For Societal Resilience. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. Article 817. doi:10.1145/3613904.3642686
- [26] Michael Oppermann, Robert Kincaid, and Tamara Munzner. 2020. VizCommender: Computing text-based similarity in visualization repositories for content-based recommendations. *IEEE Trans. on Visualization and Computer Graphics (TVCG)* 27, 2 (2020), 495–505. doi:10.1109/TVCG.2020.3030387
- [27] Michael Oppermann and Tamara Munzner. 2020. Data-first visualization design studies. In *IEEE Workshop on Evaluation and Beyond-Methodological Approaches to Visualization (BELIV)*. 74–80. doi:10.1109/BELIV51497.2020.00016
- [28] Raja Parasuraman, Thomas B. Sheridan, and Christopher D. Wickens. 2000. A model for types and levels of human interaction with automation. *IEEE Trans. on Systems, Man, and Cybernetics—Part A: Systems and Humans* 30, 3 (2000), 286–297. doi:10.1109/3468.844354
- [29] Sharon K. Parker and Gudela Grote. 2022. Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Annual Review of Organizational Psychology and Organizational Behavior* 9 (2022), 1–27. doi:10.1146/annurev-orgpsych-012420-055503

- [30] Sebastian Raisch and Sebastian Krakowski. 2021. Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review* 46, 1 (2021), 192–210. doi:10.5465/amr.2018.0072
- [31] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. 2012. Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. on Visualization and Computer Graphics (TVCG)* 18, 12 (2012), 2431–2440. doi:10.1109/TVCG.2012.213
- [32] Eva Selenko, Sarah Bankins, Mindy Shoss, Joel Warburton, and Simon Lloyd D Restubog. 2022. Artificial intelligence and the future of work: A functional-identity perspective. *Current Directions in Psychological Science* 31, 3 (2022), 272–279. doi:10.1177/09637214221091823
- [33] Leixian Shen, Enya Shen, Yuyu Luo, Xiacong Yang, Xuming Hu, Xiongshuai Zhang, Zhiwei Tai, and Jianmin Wang. 2023. Towards Natural Language Interfaces for Data Visualization: A Survey. *IEEE Trans. on Visualization and Computer Graphics (TVCG)* 29, 6 (2023), 3121–3144. doi:10.1109/TVCG.2022.3148007
- [34] Thomas B. Sheridan. 2002. *Humans and Automation: System Design and Research Issues*. Wiley. doi:10.1108/03684920310493440
- [35] Ben Shneiderman. 2022. *Human-Centered AI*. Oxford University Press.
- [36] Binyi Su, Zhong Zhou, and Haiyong Chen. 2022. PVEL-AD: A large-scale open-world dataset for photovoltaic cell anomaly detection. *IEEE Trans. on Industrial Informatics* 19, 1 (2022), 404–413. doi:10.1109/TII.2022.3162846
- [37] Susanne Tafvelin and Andreas Stenling. 2018. Development and initial validation of the need satisfaction and need support at work scales: A validity-focused approach. *Scandinavian Journ. of Work and Organizational Psychology* 3, 1 (2018). doi:10.1037/cou0000323
- [38] Michael Vössing, Niklas Kühl, Matteo Lind, and Gerhard Satzger. 2022. Designing transparency for effective human-AI collaboration. *Information Systems Frontiers* 24, 3 (2022), 877–895. doi:10.1007/s10796-022-10284-3
- [39] Magdalena Wischniewski, Nicole Krämer, and Emmanuel Müller. 2023. Measuring and understanding trust calibrations for automated systems: A survey of the state-of-the-art and future directions. In *Proc. Conf. on Human Factors in Computing Systems (CHI)*. 1–16. doi:10.1145/3544548.3581197
- [40] Heather M Wojton, Daniel Porter, Stephanie T. Lane, Chad Bieber, and Poornima Madhavan. 2020. Initial validation of the trust of automated systems test (TOAST). *Journ. of Social Psychology* 160, 6 (2020), 735–750. doi:10.1080/00224545.2020.1749020
- [41] David D Woods. 2018. The theory of graceful extensibility: basic rules that govern adaptive systems. *Environment Systems and Decisions* 38, 4 (2018), 433–457. doi:10.1007/s10669-018-9708-3
- [42] Amy Wrzesniewski and Jane E Dutton. 2001. Crafting a job: Revisioning employees as active crafters of their work. *Academy of Management Review* 26, 2 (2001), 179–201. doi:10.5465/amr.2001.4378011
- [43] Pippa Yeoman and Lucila Carvalho. 2019. Moving between material and conceptual structure: Developing a card-based method to support design for learning. *Design Studies* 64 (2019), 64–89. doi:10.1016/j.destud.2019.05.003