



AI-human learning systems: Investigating the strategic role of AI for organizational learning

Strategic Organization

2026, Vol. 24(2) 307–342

© The Author(s) 2025



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14761270251385860

journals.sagepub.com/home/soq

Patrick Figge¹, Edward Anderson²
and Kyle Lewis³

Abstract

Much attention has been devoted to the effects of artificial intelligence on individual performance for isolated tasks. To predict the indirect and longer-term effects of artificial intelligence in organizations, however, we need to better understand how artificial intelligence-enabled technology affects work that is interdependent and completed by multiple persons in collaboration. We draw from the organizational learning and strategic management literature to develop a formal mathematical model that reflects the dynamic, interdependent, non-linear relationships in the organizational systems in which artificial intelligence is increasingly embedded. Using the methodology of system dynamics, we model individual, collective, and artificial intelligence learning processes and their effects on performance over time. Our results provide insights into how artificial intelligence affects organizational learning and capabilities.

Keywords

artificial intelligence, collective knowledge, organizational learning, system dynamics, technology

Artificial intelligence (AI)-enabled technology is changing the nature of work, altering what and how employees learn and remember, and how organizations apply that knowledge to create competitive advantages. Research that informs the role of AI in organizations is rapidly increasing and has delivered important insights about the complexity of human-machine interaction (e.g. Allen and Choudhury, 2022; Bailey et al., 2022; Davenport, 2018; Raisch and Fomina, 2024), yet few studies examine how AI affects work that is interdependent and completed by multiple people in collaboration. This has led to calls for a “systems perspective” to studying AI in organizations, with human and non-human entities operating in a dynamic system of interdependent relations (Agrawal

¹University of Passau, Germany

²University of Texas at Austin, USA

³University of California, Santa Barbara, USA

Corresponding author:

Patrick Figge, University of Passau, Innstr. 27, 94032 Passau, Germany.

Email: patrick.figge@uni-passau.de

et al., 2022; Anthony et al., 2023; Bailey et al., 2022; Faraj and Leonardi, 2022). Therefore, integrating AI technologies in organizations demands a strategic effort to design organizational systems that (continue to) allow for effective collaborative relations between humans and with increasingly agentic technologies.

In this study we analyze the effects of AI that is embedded in a system of humans, groups, and the organization. We draw from the organizational learning and strategic management literature (Argote, 2013; Argote et al., 2021; Barney, 1991; Edmondson, 2002; Eisenhardt and Martin, 2000; Harvey et al., 2022; Kemp, 2024) to develop a system dynamics model of individual, collective, and AI learning processes over time, focusing on accumulated knowledge, learning, forgetting, and coordination among human and non-human entities. Our approach is consistent with calls to link organizational learning and the organizational capabilities literature (e.g. Argote, 2013; Edmondson, 2002; Harvey et al., 2022). These calls argue that connecting multiple levels of analysis (individuals, groups, organizations)¹ not only contributes to a better understanding of how AI might help organizations achieve long-term performance, but also reveals “underexplored dynamic managerial capabilities stemming from team and organizational design” (Harvey et al., 2022: 497; Rahmandad, 2012). The organizational learning lens relates to research examining the microfoundations of organizational capabilities, providing a link between decisions of senior managers and organizational capabilities (Argote and Levine, 2020).

Our “learning system” model also reveals the conditions under which intertwined learning processes at different levels of analysis can produce organizational performance. The model incorporates findings from extant empirical research, conceptual papers, and computational modeling. It allows for new insights about how interrelated learning systems of individuals, groups, and AI produce or interfere with knowledge accumulation, organizational learning, and performance. Specifically, we examine how AI-enabled technologies affect the complex processes and outcomes of work teams and larger collectives. Our research question is: How does the use of AI in a collective setting alter individual and collective learning, the accumulation of expertise, and ultimately performance? By focusing on the microprocesses that are well understood from studies of group and organization learning, we aim to shed light on how AI can best be leveraged for overall organizational performance.

We present a formal mathematical model that explains how individual and collective learning unfold when AI is embedded in a group of humans whose work is interdependent, collaborative, and depends on creating and leveraging knowledge. Knowledge-based work is common in all organizations and especially relevant for organizations whose core activities involve innovation, software or product development, data analysis, or the use and interpretation of information in general. The use of AI is becoming normative in such organizations, which are leveraging the power of AI to try to increase human performance (Agrawal et al., 2022; Anthony, 2021; Dell’Acqua et al., 2023). Despite the increasing use of AI, we still know little about the effects of using AI in collective work processes that up until this point have been quite well understood (Argote, 2013; Argote et al., 2021; Argote and Miron-Spektor, 2011). Moreover, much of the current research on AI examines AI as isolated and independent from the rest of the organization (Agrawal et al., 2021, 2022; Raisch and Fomina, 2024). For these reasons, several authors have argued for a systems approach to studying AI in organizations (Agrawal et al., 2022; Anthony et al., 2023; Bailey et al., 2022; Faraj and Leonardi, 2022).

Our mathematical model reflects the dynamic, interdependent, non-linear relationships of systems in which AI is embedded. Complex interdependencies and non-linear relationships that evolve over time can be modeled with system dynamics (Sterman, 2001, 2010), yielding insights that are difficult to observe through traditional empirical methods. The model’s assumptions and parameters are based on established theory in the organization and strategy science literatures and

extant empirical research on AI. Based on this model, we conduct computational experiments to compare different ways AI influences individual and collective learning.

Background and theory

AI-enabled technology

Emerging technologies are equipped with AI features, participating in organizational processes and performing cognitive tasks that previously required human knowledge and capabilities. Increasingly “AI or cognitive technologies employ such capabilities—previously possessed only by humans—as knowledge, insight, and perception to solve . . . tasks.” (Davenport, 2018: 9). These technologies are (to varying degrees) autonomous and can interact with collaborating humans in teams (Hu et al., 2021; O’Neill et al., 2022, 2023; Raisch and Fomina, 2024) and have the ability to learn from data (e.g. Wang et al., 2023), either independently or with human intervention. Fiore and Wiltshire (2016) even stipulate that technologies should be designed to include “scaffolding” functionality that “directly supports team-level processes by helping to mediate and support the interaction between individual and team-level cognitive activity” (p. 11). We refer to AI-enabled technology as “AI” throughout this article.

While AI is evolving from a tool to becoming a counterpart in teamwork (Anthony et al., 2023; Malone, 2018), prior research predominantly analyzes AI and their “users” in isolation from organizational processes and at the level of individual tasks. Recent attempts to build a theoretical foundation for inquiries into collaboration with AI have repeatedly used classic frameworks (such as input-process/mediator-output) from team research (O’Neill et al., 2022; Sebo et al., 2020; You and Robert, 2017). Research has mostly focused on the effects of AI use by individuals, but this is changing. Some studies examine how AI-enabled technologies (especially robots) affect teams and complex social settings (Jung and Hinds, 2018; Sebo et al., 2020). Findings from these studies suggest that not only are AI usage patterns markedly different in teams compared with usage by individuals (Sebo et al., 2020), but also that AI affects the way that team members interact *with each other*. For instance, Traeger et al. (2020) show that AI communication (with or without empathy) affects how team members communicate with each other. In their review of 101 research articles, Sebo et al. (2020) distill from prior research that verbal and non-verbal behavior of embodied technologies (e.g. eye contact, gestures) can shape human perception and anchor the way human members interact with each other.

Despite these advances, we still have a very limited understanding of how team learning, knowledge, and social dynamics are affected when teams work together with AI (Argote et al., 2021; Jung and Hinds, 2018). Answers to the question of how humans and AI can work together often focus on how AI can be designed, rather than on how collaborating human workers need to adapt, learn, organize, and interact with AI for high performance (Wang et al., 2023). Team learning is emergent: it “originates in the cognition, affect, behaviors, or other characteristics of individuals, is amplified by their interactions, and manifests as a higher-level, collective phenomenon” (Kozlowski and Klein, 2000: 55). These complex interdependencies are difficult to study with traditional empirical techniques.

Some ethnographic research does, however, offer insights into the complex adaptive dynamics involved in adopting AI-enabled technology. For example, Beane (2019) found that the use of surgical robots altered the practices of operating teams in a way that rendered prior learning strategies of surgical trainees less effective. Barrett et al. (2012) found that teams needed to re-negotiate role structures when the relative skills or status between occupational groups changed as a result of integrating a programmable dispensing robot into a hospital’s systems. Beane and Orlikowski

(2015) and Sergeeva et al. (2020) showed that the knowledge that team members need to access from each other and the accompanying coordinative processes were different when using AI-enabled technologies. These technologies change how humans interact (via robotic telepresence, see Beane and Orlikowski, 2015) or what they can do (enabled by a surgical robot, see Sergeeva et al., 2020). Informed by these deep descriptions of the intricate changes in teams that work alongside AI, we aim to conceptualize how individual and collective learning change when teams collaborate with AI.

When AI is embedded in organizational learning systems, AI is capable of influencing the learning and knowledge of human actors and is itself a learning entity in the system (Anthony et al., 2023; Kemp, 2024). AI-enabled technologies are able to “continuously acquire knowledge and skills, possibly operating autonomously or in concert with humans” (Bailey et al., 2022: 2). Understanding these complex systems of human and non-human entities requires researchers to “conceive of technologies as constituted by relations and interleaved with other relations that are always evolving” (Bailey et al., 2022; Faraj and Leonardi, 2022: 777). Too often, however, AI and algorithms are treated as a black box, and different types of organizations are treated similarly (Zeng and Glaister, 2018).

Understanding the dynamic and long-term effects of using AI-enabled technologies in organizations would go a long way toward explaining why some organizations perform better when using the same technology (Kemp, 2024; Zeng and Glaister, 2018). While much evidence points toward AI augmenting individual capabilities for many tasks (e.g. Dell’Acqua et al., 2023; Noy and Zhang, 2023), we do not yet understand how AI changes the complex relationships that comprise teams and organizations. Past research has been conducted mostly within disciplinary silos, and, in particular, outside of strategy, organization science, and organizational learning research. Beyond questions on how to design technology to make it more suited for humans to work with, we need a more comprehensive strategic understanding of how AI technologies can be leveraged by organizations. In this article, we offer a new, strategic organization perspective informed by organizational learning theory to analyze how organizations can redesign their work systems to support learning, collaboration, and performance.²

Organizational learning and forgetting

Organizational learning theory is especially suitable for explaining organizational processes that span different levels of analysis, such as individuals, groups, and organizations (Argote, 2013; Argote et al., 2021; Crossan et al., 2021; Huber, 1991; Shapira, 2020). As we explain in this article, organizational learning theory can also explain the relationships between AI, human learning, and organizational knowledge (Argote et al., 2021; Sturm et al., 2021).

Learning from experience, by individuals, groups, or organizations, is often operationalized as learning curves (also known as experience curves, progress curves, and learning by doing) (Argote, 2013). A learning curve represents the relationship between cumulative experience and performance improvements. Since Wright (1936) documented the industry learning curve, hundreds of studies have examined learning curves for a variety of organizational settings, outcomes, and across different levels of analysis (e.g. Anderson and Parker, 2002; Argote, 1993; Argote and Epple, 1990; Nembhard and Uzumeri, 2000; Reagans et al., 2005; Sterman et al., 1997).

Research on organizational learning has revealed various mechanisms through which organizations can learn. One prominent mechanism is learning by doing, where organizations acquire knowledge through direct experience and experimentation. This can involve trial and error, experimentation with new processes or technologies, and reflecting on the outcomes. Another mechanism is learning from mistakes. By analyzing past failures and understanding their underlying

causes, organizations can identify areas for improvement and develop strategies to prevent similar mistakes in the future (see Argote et al., 2021; Argote and Miron-Spektor, 2011 for literature reviews on organizational learning).

Retaining knowledge accumulated via organizational learning is inherently imperfect. Retained knowledge may become obsolete because of external changes (e.g. industry changes or technological advancements) or decisions by the organization (e.g. a new strategy that requires different knowledge). Knowledge may also be forgotten, for instance, when it remains unused for too long or when the individuals who remember it leave the organization (and the knowledge was not codified). Organizational forgetting refers to this phenomenon, where an organization's knowledge depreciates in value or is lost over time.

An important insight of the organizational learning literature is that knowledge gained from experience becomes embedded not only in the minds of humans, but also in routines, processes, structures, and tools of the organization (Argote, 2013; Argote et al., 2021). "In order for learning to be organizational, the knowledge an individual acquires would have to be embedded in a supra-individual repository . . . so that the knowledge would persist in the organization even if the individual were to depart" (Argote et al., 2021: 5403). One such supra-individual repository is a *transactive memory system* (TMS), a collective system for encoding, storing, and retrieving information that develops as group members gain experience working together (Lewis and Herndon, 2011; Ren and Argote, 2011). Research on TMS in groups shows that developing a TMS is strongly associated with higher group performance for a variety of tasks and contexts (Bachrach et al., 2019). This is attributed to several factors, including accurate task assignments, fast knowledge access, and progressively specialized individual knowledge (Lee et al., 2014; Lewis, 2003, 2004). Developing a TMS is especially advantageous for knowledge worker teams, whose performance depends on developing and leveraging expertise.

Anderson and Lewis (2014) modeled the learning processes involved in developing a TMS. Two key findings from their study are worth mentioning here. First, when individuals could rely on their shared (collective) knowledge of "who knows what" in the group, they were more apt to develop specialized expertise, refine their collective understandings of "who knows what" and further accelerate individual learning. These cycles of learning at different levels of analysis (Lewis et al., 2005) are positively reinforcing over time, resulting in increased performance. Second, members' knowledge contributes to productivity only up to a point—if members' knowledge becomes "overspecialized," the group's performance declines. Without enough shared (collective) knowledge, highly specialized members may have difficulty finding common ground (e.g. Clark, 1985, 1998; Krauss and Fussell, 1996), as specialized vocabulary, references, and knowledge may not be understood by members with disparate expert knowledge (Fraidin, 2004). Without sufficient collective knowledge, members will be unable to build on and integrate other members' disparate knowledge, which in turn further slows collective learning.

Thus, *both* individual and collective knowledge are needed for high performance on tasks completed interdependently by multiple individuals. The research on AI and teams rarely considers effects of collective learning, nor its role in facilitating individual learning, despite calls to study AI in the context of interdependent action (e.g. McNeese et al., 2023; O'Neill et al., 2022). As Anderson and Lewis (2014) showed, individual and collective learning processes are intertwined and dynamic, with learning at the individual level reciprocally linked to learning at the collective level. Anderson and Lewis (2014) described these interrelated learning curves as a *learning system*. We adopt the same terminology to describe a larger learning system that incorporates AI learning into the system of individual and collective learning: the AI-Human Learning System.

We construct a computational model of the AI-Human Learning System consistent with the learning curves literature, incorporating learning rates, forgetting rates, and intertwined learning

processes that connect individual learning with collective learning (group, unit, or organization). We consider AI itself to be a learning entity that accumulates knowledge as a function of task completion. AI learning can be construed as refinements or expansions in AI algorithms that tend to increase overall performance in a given organizational context. AI also learns by extracting patterns from a large number of cases or task iterations.³

In the next section, we describe a computational model, conceived as a learning system that combines AI learning, individual learning, and collective learning. Our model recognizes that AI is embedded in a system of interactions that unfold over time, and which include both human and non-human actors (Anthony et al., 2023; Bailey et al., 2022; Faraj and Leonardi, 2022). We run computational experiments with the model, systematically varying the potential influences of AI on individual and collective learning. Thus, our model does not emphasize specific features of AI but instead emphasizes the impact of AI on accumulated knowledge and learning at multiple levels of analysis. By systematically varying model parameters, we create different cases (conditions) wherein AI augments individual or collective learning or causes obsolescence of individual or collective knowledge. Comparing these different cases provides insights that are applicable for a variety of knowledge-intensive organizations and settings.

AI-human learning system

Systems are a set of relationships between actors and related interdependent components that work together to perform whatever functions are required to achieve the system's objective (Meadows, 2008). The learning system model presented in this article assumes that organizational outcomes are a function of learning from multiple human and non-human entities that work together to complete tasks. The combined effects of learning from AI and interacting humans are depicted in the *causal loop diagram* in Figure 1. The variables are consistent with learning curve research and include knowledge stocks, rates of learning and forgetting, and the rate of task completion (Anderson and Lewis, 2014; Anderson and Parker, 2002). Stocks are accumulations of knowledge within the system. In our case, the learning system is represented by three knowledge stocks that influence the system's performance (its task completion rate). Each of the three knowledge stocks is attached to a learning rate (inflow) and a forgetting rate (outflow). Inflows increase knowledge stocks and outflows decrease knowledge stocks. The causal loop diagram is a stylized and simplified version of the formal mathematical model that helps reveal the key variables and relationships in the learning system.

We model AI itself as a learning entity, such that AI accumulates knowledge with each task that is completed by the group, in the given context. We assume that AI learning is intertwined with the individual and collective learning systems, such that the AI learning system affects the other learning systems, and vice versa. As stated in Sturm et al. (2021), "humans are no longer the only ones capable of learning and contributing to an organization's stock of knowledge" (p. 1581).

A few assumptions undergird our model. First, we assume that the task is one that AI has the capacity to address. Second, we assume that the task is one that AI cannot accomplish by itself—that is, both human and non-human entities need to coordinate for high task performance (cf. Raisch and Fomina, 2024: 4–5). This is captured by the way that performance is calculated: as a function of the contributions from the individual, collective and AI knowledge stocks.

Mathematical specification

In this section, we describe the mathematical specification of the AI-Human Learning System.⁴ The simulation model described by the equations below uses the Stella v. 2.1.2 simulation




Figure 1. Causal loop diagram of the AI-human learning system.

Notes. Solid arrows indicate a positive directionality of a causal relationship. For example, the relationship between the individual contribution to productivity and task completion rate is positive, indicating that as the individual contribution increases, so does the task completion rate. A dashed line indicates a negative polarity: when the preceding construct increases, the construct on the arrowhead decreases. The word “delay” on an arrow indicates a time delay in the effect. For example, the relationship between the individual learning rate and individual knowledge is not instantaneous: it takes time for an individual to learn and develop new knowledge. These delays lend validity to the model, as most learning effects emerge over time; they are not instantaneous. Three knowledge stocks are represented in the diagram: individual knowledge, collective knowledge, and AI knowledge. The task completion rate is a function of (is linked to) all three of these knowledge stocks. Also represented in the diagram is the potential for overspecialization of individuals, which limits communication and coordination and reduces the collective learning rate. It is the collective learning rate that increases the collective knowledge stock, which allows individuals to further specialize their expertise. The drawing also incorporates forgetting, for each of the learning entities.

application. The model is available upon request from the authors. The simulation executes every time period t , which we interpret as 1 week. At that point, all of the variables are recalculated. Many of the variables and parameters related to collective, individual, and AI knowledge are analogous, and so will be represented by the same variable name with a superscript of “ c ,” “ i ,” and “ a ,” respectively. Hence, K_t^c represents the collective knowledge stock at time t , and K_t^i the individual knowledge stock at time t . That said, the resultant learning curves for individuals and the collective will not be identical, because these different entities tend to accumulate different types of knowledge at different rates (Reagans et al., 2005).

Performance. Performance q_t is computed at each time period, where t is the time period in weeks. Performance is calculated by multiplying the contributions from the collective (c_t), individuals (i_t), and AI (a_t) to performance, or

$$q_t = c_t i_t a_t \quad (1)$$

Learning and forgetting. Performance depends not only on knowledge accumulated through learning, but also on knowledge lost to “forgetting” (Argote, 1993, 2013; Argote et al., 2021; Wright, 1936). Forgetting occurs when knowledge becomes obsolete because of new or changing task conditions, when knowledge leaves the group via turnover, or when AI takes over aspects of the task previously completed by humans. Learning results from the tasks completed in prior time periods, and forgetting results from the depreciation of knowledge accumulated in prior periods (Darr et al., 1995).⁵ We begin with the specification for individual knowledge that captures these effects, following Anderson and Lewis (2014):⁶

$$K_{t+1}^i = K_t^i + \lambda_i l_t^i q_t - \chi^i K_t^i \text{ where } K_0^i = 1 \quad (2)$$

K_t^i is the accumulation or stock of individual knowledge at time t . Individual learning, which represents the accumulation of learning per task completed, is the product of λ_i and l_t^i . The parameter λ_i is a constant that represents the base rate of learning per task completed. In contrast, the value of l_t^i can vary during a simulation run to represent the impact on individual learning by specialization. The parameter χ^i is the individual knowledge fractional forgetting rate, which is a constant for any given simulation unless explicitly stated otherwise. It represents effects such as individual turnover, cognitive limitations, or knowledge becoming obsolete over time (Argote, 2013; Argote et al., 2021).

As with individual knowledge, collective (K_t^c) and AI (K_t^a) entities accumulate and forget knowledge, so their equations and parameters are analogous:

$$K_{t+1}^c = K_t^c + \lambda_c l_t^c q_t - \chi^c K_t^c \text{ where } K_0^c = 1$$

and

$$K_{t+1}^a = K_t^a + \lambda_a l_t^a q_t - \chi^a K_t^a \text{ where } K_0^a = 1 \quad (3)$$

For convenience and without loss of generality, all three knowledge stocks initially have a value of 1. Overspecialization can change the value of l_t^c during the simulation dynamically. Individual learning and forgetting, as well as collective learning and forgetting, can also be influenced by AI and how it is used in the organization, which we will examine later by varying λ_i , λ_c , χ^i , and χ^c in the computational experiments.

Finally, it should also be noted that K_t^i is a single stock that represents the average of each individual’s knowledge. Similarly, K_t^a is a single stock that can represent the aggregate knowledge of more than one AI when more than one AI application is integrated in the team.

Learning curve for individuals. The contribution of each knowledge stock to performance is the standard power-law formulation. For individual knowledge,

$$i_t = \left(\frac{K_t^i}{K_0^i} \right)^{\gamma_i} \quad (4)$$

γ_i represents the individual learning curve parameter. It is very rare for learning curve parameters to fall outside the range of 0.0 and 1.0, so we restrict it to these values. In practice, it seems to most commonly fall between 0.1 and 0.5 (Argote, 2013). This causes performance to increase at a diminishing rate with increased knowledge.

Learning curve for collective and AI knowledge. The contributions of the collective and AI knowledge stocks are calculated analogously with their respective constants γ_c and γ_a representing the learning parameter for accumulating collective and AI learning:

$$c_t = \left(\frac{K_t^c}{K_0^c} \right)^{\gamma_c} \text{ and } a_t = \left(\frac{K_t^a}{K_0^a} \right)^{\gamma_a} \quad (5)$$

Opportunity to specialize. When members learn together, new collective knowledge (about who knows what, shared vocabulary, shared routines) can develop, and an adaptive cycle of learning can unfold (Anderson and Lewis, 2014). As collective knowledge increases, it affects what, and how quickly, individuals learn. TMS research suggests that when individual members can rely on collective knowledge, it provides cognitive slack for deeper individual learning—typically in knowledge areas that are not already covered by other members (Hollingshead, 2001; Lewis et al., 2007; Wegner et al., 1985; Wittenbaum et al., 1996). This “progressive specialization” can benefit performance because it allows a greater amount of task-relevant knowledge to be applied to the group’s tasks. As collective knowledge increases, individuals can increase their expertise more effectively, albeit with diminishing returns. We specify this effect based on the standard literature of diminishing returns to functions of knowledge or other attributes (Ellis, 1965):

$$l_t^i = \left(\frac{K_t^c}{K^s} \right)^{\sigma} \quad (6)$$

σ is a parameter that determines the increase in learning. It is restricted to be less than 1 to ensure diminishing returns to knowledge. K^s is another normalization constant, such that when K_t^c reaches K^s , the effect of specialization on individual learning has a value of 1, which intuitively is the point at which the specialization effect begins to become less important.

Ability to coordinate. If members’ knowledge becomes too specialized, it can impede communication and reduce members’ ability to build on and integrate knowledge from other members of the group. The result is a slowing of collective learning and decreased collective knowledge. Because learning entities are not independent but are intertwined, we can expect that lower collective knowledge will eventually slow individual learning. The interplay between individual knowledge, collective knowledge, and subsequent learning can be represented mathematically with a reverse “S-curve.” S-curves are often used to represent the effects of parameters in learning models (Repenning, 2002; Sterman et al., 1997). Specifically, we use the formulation of the S-Curve by Repenning (2002) as adapted by Anderson and Lewis (2014) and specify the curve as

$$I_t^c = 1 - \left[1 + \exp \left(-\omega \left[\frac{K_t^i}{K^o} - \frac{1}{2} \right] \right) \right]^{-1} \quad (7)$$

where K^o is a normalization constant so that the reduction in the ability to coordinate reduces collective learning by half. Intuitively, it can be thought of as the point where overspecialization “kicks in.” ω is a parameter that determines the inverse S-curve’s maximum slope at its midpoint.

Model testing and simulations

We conducted standard simulation robustness tests appropriate to system dynamics-based models of theory (Sterman, 2010). Table 1 shows the parameters we use for a “base” case, which we later compare with simulations from computational experiments. We use the parameter values that were validated by Anderson and Lewis (2014); these same parameter values are informed by empirical research on learning and forgetting (e.g. Argote, 2013; Boone et al., 2008; Darr et al., 1995). Unless stated otherwise, all the variables and parameters will always be greater than or equal to zero. We also conduct sensitivity tests (Monte Carlo analyses, see Appendix 2) that vary parameter values to ensure that our model, which adds AI as a learning entity, is robust and that the results of our computational experiments are valid for a broad range of parameters.

Base Case simulation. Simulations using the mathematical specifications above are illustrated with Figure 2. Note that the x-axis is time, and the y-axis represents performance. Each “trace” (or curved line) represents the behavior of the model over time.

Results from the Base Case simulations show that performance of the AI-Human Learning System increases up to week 25. The peak performance of the AI-Human Learning System is 12.5 tasks per week, which is higher than the peak performance of a learning system model *without* AI.⁷ Thus, we see a performance increase when AI is integrated with human learning activities. After peak performance is reached, however, performance begins to decline, until about week 168. The reasons for this decline are explained below. AI use appears to delay this decline, but it does not eliminate it. At week 168, the task completion rate begins to oscillate as it increases—albeit slightly—again. The reason is that eventually the task completion rate drops so low that individuals become less specialized, eventually allowing collective learning to build up again, kicking off a weak virtuous cycle that peaks at week 195.

Interpreting the Base Case simulation results. To get a deeper understanding of these Base Case results, we disentangle effects emanating from the different learning entities (collective, individual, and AI). The results are shown in Figure 3. We “zoom in” on the weeks 0–120 for convenience. Figure 3 shows each of the knowledge stocks (collective, individual, and AI) over time, and their relative contribution to performance.

Initially, collective, individual, and AI knowledge all increase, driving up each of their contributions to performance (bottom right). Individual knowledge and its contribution to performance is accelerated from the beginning by increased collective knowledge, which allows individuals to develop new specialized knowledge. As members’ knowledge becomes overspecialized, however, communication and coordination eventually become less effective, progressively reducing the accumulation of collective knowledge. When the rate of collective learning falls below collective forgetting at week 5, collective knowledge and its contribution to performance begin to decline. In

Table 1. Base model parameters.

Parameter	Meaning	Value	Units
K_0^i	Initial stock of individual knowledge	1	Individual knowledge units
K_0^c	Initial stock of collective knowledge	1	Collective knowledge units
K_0^a	Initial stock of AI knowledge	1	AI knowledge units
χ^i	Individual fractional forgetting rate	1/24	1/weeks
χ^c	Collective fractional forgetting rate	1/36	1/weeks
χ^a	AI fractional forgetting rate	1/24	1/weeks
γ_i	Individual learning curve parameter	0.2	Dimensionless
γ_c	Collective learning curve parameter	0.4	Dimensionless
γ_a	AI learning curve parameter	0.2	Dimensionless
K^0	Normalization constant for overspecialization	10	Individual knowledge units
ω	Overspecialization parameter	5	Dimensionless
K^s	Normalization constant for specialization	5	Collective knowledge units
σ	Specialization parameter	0.4	Dimensionless
λ_i	Base individual learning rate per task completed	1	Individual knowledge units per task completed
λ_c	Base collective learning rate per task completed	1	Collective knowledge units per task completed
λ_a	Base AI learning rate per task completed	1	AI knowledge units per task completed

addition, individual and AI contributions to performance begin to flatten from diminishing returns to knowledge. The result is that collective knowledge’s contribution begins dragging down increases to performance, until at week 25 the loss overtakes the flattening contributions from individual and AI knowledge, causing overall performance to drop. At this point a vicious cycle begins. The performance decline reduces learning, which in turn slows the accumulation of individual and AI knowledge. Eventually, the individual learning rate drops below that of the individual forgetting rate, which results in a decline of individual knowledge and its contribution to performance at time 45. The same happens to AI learning, knowledge, and its contribution to performance. Since the creation of new AI knowledge slows down with the dwindling performance, AI’s learning rate drops below its forgetting rate as well.

This “overshoot” behavior is an established pattern in system dynamics research, reflecting the dynamic, non-linear, and interdependent relationships in the model and, specifically, the interrelated learning curves that capture organizational learning. Prior research in system dynamics has uncovered such “overshoot” patterns time and again when modeling learning or process improvement activities using learning curves (e.g. Anderson and Chandrasekaran, 2024; Anderson and Parker, 2002; Repenning and Sterman, 2001; Sterman et al., 1997). Anderson and Lewis (2014) added “overspecialization” as a reason for why this pattern occurs to the literature on organizational learning systems.

Interpreting the impact of AI on human learning. In the Base Case, performance declines over the long run as the team loses its capability to use the AI and integrate it appropriately in their




Figure 2. Base Case.
Notes. The simulation of the Base Case uses the parameter values shown in Table 1. During the 240-week simulation, the Base Case has the following summary statistics: Average=5.3; Maximum=12.5; Final value=2.1 [units: tasks completed per week].

collaborative work. This, combined with overspecialization of members’ knowledge, creates a performance decline that is difficult to avoid or repair.

While the effects of AI use on the human learning system may seem counterintuitive, there is new research that explains why using AI in collaborative environments can be detrimental in the long term. Specifically, negative long-term effects of AI use have been documented in case studies of AI coding assistants (Anderson et al., in press). In controlled environments on isolated tasks, there is usually a clear performance benefit for individual human coders using AI coding assistants. However, AI-generated code in real organizational (legacy) environments may create “technical debt,” that is, the need for future rework due to, for instance, an entangled web of dependencies that make it difficult for human coders to contribute meaningfully in the future (Anderson et al., in press). In addition, when AI is used by individual coders that are working as a team, integrating that code takes time away from production activities. Informants in informal interviews conducted for this study noted that it was hard to keep track of “who knows what” after members individually used AI coding tools. For tasks that require the integration of human and AI knowledge, this can be detrimental for task performance in the long run.

Computational experiments. The Base Case simulation showed that AI can amplify early performance by increasing a team’s capabilities to carry out tasks. To further explore how AI affects the human learning system, we conduct computational experiments that systematically vary model parameters to produce four cases that can be compared with the Base Case and with each other. The cases are shown below in Table 2 and explained below.

Recall that performance, in our model, is a multiplicative function of collective, individual, and AI knowledge stocks. Focusing on knowledge stocks, we capture an AI’s influence in the AI-Human Learning System by varying the rates at which the knowledge stocks are filled by learning and depleted by forgetting. We consider *all* inflows and outflows into the knowledge stocks that represent human learning. When AI increases a learning rate, it results in a faster *accumulation* of knowledge in the knowledge stocks for individuals and the collective, respectively. When AI




Figure 3. Base Case details: knowledge stocks and their contribution to performance.

Notes. The simulation of the Base Case uses the parameter values shown in Table 1. To make the figures readable and easier to interpret, we zoom in on the first 120 weeks, that is, we only show the first half of the 240 weeks simulation run that we show in Figure 2. In the last graph, multiplying the contributions to task completion made by the three knowledge types yields the task completion per week that we show in Figure 2. For instance, in week 60 the task completion rate is 9.49 tasks/week (see Figure 2). The values in the last graph in this Figure are 1.13 for the contribution by Collective Knowledge, 2.81 for Individual Knowledge, and 2.99 for AI Knowledge. Multiplying these three values yields 9.49.

increases a forgetting rate, it results in a faster *reduction* in the knowledge stocks for individuals and the collective, respectively. In the next sections we provide examples that illustrate each case, results of model simulations, and interpretations of those results.

Table 2 shows four different cases representing different scenarios for how AI could impact the AI-Human Learning System. For example, AI can accelerate the learning rate for individuals (first column of the table) or the collective (first row of the table), augmenting the knowledge stocks that drive performance. Alternatively, AI may cause individual or collective knowledge to become obsolete. This is modeled as increases in the forgetting rates for individuals (second column) and the collective (second row), leading to a reduction in the knowledge stocks that drive performance.

Table 2 also explains in practical terms how varying learning and forgetting rates can manifest in the AI-Human Learning System. For example, when AI increases the individual learning rate, it

Table 2. Computational experiments of AI’s influence on the AI-human learning system.

		AI’s Influence on Individuals	
		Increased Learning Rate λ_i AI augments human learning	Increased Forgetting Rate χ^i AI causes obsolescence of individual knowledge
AI’s Influence on the Collective	Increased Learning Rate λ_c AI manages role structure	Case #1 AI use increases individual knowledge or skill AI monitors, tracks, analyzes division of cognitive labor	Case #2 AI use depreciates individual knowledge, makes it obsolete, or takes on tasks previously performed by members AI monitors, tracks, analyzes division of cognitive labor
	Increased Forgetting Rate χ^c AI causes obsolescence of collective knowledge	Case #3 AI use increases individual knowledge or skill AI disrupts role structure	Case #4 AI use depreciates individual knowledge, makes it obsolete, or takes on tasks previously performed by members AI disrupts role structure

can be said that AI *augments* human learning. When AI increases the individual forgetting rate, it means that AI causes obsolescence of individual human knowledge. Practically, this type of forgetting can occur if AI takes on tasks previously performed by members, making members’ knowledge less relevant for a task. When AI increases the collective learning rate, it can be a result of AI taking on team “management” activities such as monitoring, tracking, and analyzing the division of cognitive labor in the team (Fiore and Wiltshire, 2016). When AI increases the collective forgetting rate, it may indicate that AI has disrupted the existing role structure that would otherwise help members understand “who knows what” and “who will do what.” Each case reflects effects that are emerging in the AI literature, which we describe next.

Case #1. First, we consider a case where AI increases both the collective learning rate λ_c and the individual learning rate λ_i . In Case #1, using AI increases individual skill and/or deepens individuals’ specialized expertise, while “managing” the division of cognitive labor in the group. Organizations increasingly use algorithmic management practices (e.g. Duggan et al., 2020; Kellogg et al., 2020; Parent-Rochelleau and Parker, 2022; Schweitzer and De Cremer, 2024), defined “as a system of control where self-learning algorithms are given the responsibility for making and executing decisions affecting labour, thereby limiting human involvement and oversight of the labour process” (Duggan et al., 2020: 119). In these examples, AI is applied to advise managers or to take over managerial roles that include assigning tasks, creating work schedules, and monitor task execution (Parent-Rochelleau and Parker, 2022; Schweitzer and De Cremer, 2024).

Experiments show that it is indeed possible for a technology to shape human behavior via certain “nudges” and influence a group’s collective intelligence (Gupta et al., 2024). Using AI to monitor, track, and analyze the role structure could mitigate problems of discoordination that occur

when members become overspecialized. Recent technical developments and theoretical conceptualizations point toward a path where future technologies actively influence individual and collective learning in an AI-Human Learning System.

Figure 4 shows the dynamic effects of Case #1, compared with the Base Case described earlier. When AI increases both the collective and individual learning rates, performance is higher than in the Base Case throughout most of the simulated timeframe, but it does not prevent a decline beginning in week 28. Like in the Base Case, the decline is ultimately a result of increasing specialization of the individuals' expertise, which hampers communication and coordination and slows the growth of collective (shared) knowledge. In the short term, the faster collective learning rate in Case #1 enables the team to better cope with overspecialization. Eventually, however, it accelerates the rate of tasks completed per week, *indirectly* raising the amount of specialized individual knowledge that accumulates from task experience. Hence, overspecialization will over time still slow collective learning and will eventually cause a decline in performance. That said, Case #1 shows that AI that enhances both collective and individual learning can enable the organization to reach greater levels of performance versus the Base Case throughout most of the simulation even after the downturn begins. The only exception is after week 173, Case #1's performance falls below the Base Case, but that said, it quickly recovers, finishing on par with the Base Case at the end of the simulation.

Case #2. Next, we consider a case in which AI increases the collective learning rate λ_c but, at the same time, increases the individual forgetting rate χ^i . That is, AI causes obsolescence of individual knowledge while also supporting the group's role structure. There are existing AI technologies that are representative of this case. For instance, AI agents like Microsoft's Cortana are able to "facilitate communications between people or on behalf of people, such as by transcribing a meeting and distributing a voice-searchable version to those who couldn't attend" (Wilson and Daugherty, 2018: 6), reinforcing some collective knowledge about "who knows what." The dissemination and the capabilities of such agents have recently gotten a boost thanks to advances in text processing and generation based on large language models and generative AI. Such agents can also be designed to autonomously complete (parts of) a task by themselves and/or decisively change the knowledge that is needed for humans to complete a task, thereby replacing the need for humans to have that specific knowledge. Moreover, individuals' knowledge becomes outdated (faster) when AI takes over additional parts of a task. For instance, strategy consultants have been shown to increase their performance in solving a typical and realistic problem with generative AI agents, with top-performers prior to the introduction of the AI agents profiting less (Dell'Acqua et al., 2023). In that example, individual knowledge has become obsolete to a certain extent. Similarly, AI agents can write and suggest code, enabling higher performance and proficiency in more programming languages. This allows software developers to increase the number of completed tasks, and, at the same time, with less experienced developers profiting more (Cui et al., 2024).

Figure 5 shows the dynamic effects of Case #2, compared with the Base Case. Initially, this case behaves similar to the previous case (Case #1, see Figure 4) with a similar peak magnitude followed by a decline. The computational experiment that created this case shares an increased collective learning rate with Case #1. However, this computational experiment also simulates an increase in the individual *forgetting* rate whereas in Case #1 increased the individual *learning* rate. Despite the increased individual forgetting rate, faster collective learning allows performance to rapidly increase. Eventually, greater task experience increases overspecialized individual knowledge, resulting in a long decline in performance after week 24.

In the long run, Case #2 is superior to Case #1. This is somewhat counterintuitive at first sight. The key to explaining why an increase in individual forgetting rather than an increase of individual learning seems beneficial is that there are diminishing returns to performance for both collective




Figure 4. Case #1 compared to Base Case.

Notes. To create the two traces shown in Figure 4, we varied two constants. The individual learning rate λ_i is 1.0 unit of knowledge per completed task in the Base Case and 2.0 units of knowledge unit per task completed in the Case #1 simulation. The collective learning rate λ_c is 1.0 knowledge unit per task completed in the Base Case and 2.0 knowledge units per task completed in the Case #1 simulation. During the 240-week simulation, Case #1 has the following summary statistics: Average = 6.4; Maximum = 16.1; Final value = 1.8 [units: tasks completed per week].

and individual knowledge. In the Case #1 simulation, there is—in the long run—a relatively high level of individual knowledge and a relatively low level of collective knowledge. The learning system is “stuck” during the rest of the simulation period; it does not recover from this situation and remains at a relatively low performance level. In the Case #2 simulation, however, individual learning is kept to a more moderate level because the individual forgetting rate is increased as part of the computational experiment. This constrains the impact of overspecialization, permitting collective learning to increase to a higher level. The increase of the collective contribution to performance more than makes up for a lower individual contribution to performance resulting from less specialized personnel. The result is that performance stabilizes at a significantly earlier point in time (around week 150) and at a higher average level than the performance level at the end of Case #1.

Case #3. This case represents a situation in which AI augments human skill or specialized expertise, but at the same time disrupts the group’s role structure. The effects of AI on collective processes have been examined mostly in ethnographic studies, which show how the introduction of technology into a team may lead to changes in the team’s role structures, on which employees depend to effectively work together. Adapting to the technology often means redefining the role structure and creating the need for different human expertise, that is, to learn new things both individually and collectively. For example, several studies showed that integrating intelligent robots into medical teams prompted a significant reconfiguration in existing team roles, responsibilities, and coordination processes (Barrett et al., 2012; Sergeeva et al., 2020). Such changes in the role structure also had longer-term effects on individual expertise. For example, deploying robots in surgical teams enhanced aspects of the surgeons’ access to information, perception, or operational dexterity (Beane, 2019; Sergeeva et al., 2020). Individual expertise was developed in one group (surgeons using the robot), and at the same time, nurses and anesthesiologists adapted to new roles that required them to learn new specialized knowledge (Sergeeva et al., 2020). An




Figure 5. Case #2 compared to Base Case.

Notes. To create the two traces shown in Figure 5, we varied two constants. The individual forgetting rate χ^i is 1/24 or 4.17% per week in the Base Case and double (8.33% per week) in the Case #2 simulation. The collective learning rate λ_c is 1.0 unit of knowledge per completed task in the Base Case and 2.0 in the Case #2 simulation. During the 240-week simulation, Case #2 has the following summary statistics: Average = 7.0; Maximum = 17.1; Final value = 3.4 [units: tasks completed per week].

organizational learning lens helps us understand the short- and long-term effects of such role structure (understanding “who knows what” and “who does what”) change and the development of new individual expertise, relevant for a new role structure.

An example for this might be the drug discovery with machine learning case that Sturm et al. (2021: 1585) describe. Scientists are provided with novel insights on molecule substructures and cause-effect relationships of how these structures contribute to pharmacological effects. Based on these insights, they may accumulate individual knowledge faster. However, the human research teams still retain agency over the process, deciding on which learning algorithms to use, setting up data for the learning algorithms, deciding on the direction that their research will be taking (e.g. which disease will be targeted, which specific process that leads to a pathogen’s effect will be targeted). The advantage of these machine learning for drug discovery techniques explicitly lies in detecting “previously disregarded substructures in molecules” and machine learning systems may “not necessarily follow the same reasoning that is provided by existing textbook knowledge of chemistry,” (Sturm et al., 2021: 1585). Those new insights may help scientists to advance more quickly but also introduce new roles to the process (e.g. data scientists) and may require new—even surprising—changes to the collaborative team efforts and to co-specialization requirements.

Figure 6 compares the simulation of Case #3 with the Base Case. Because AI increases the collective forgetting rate, the accumulation of collective knowledge is slowed down. It is more difficult to accumulate enough collective knowledge before overspecialization sets in. Hence, the organization cannot allocate roles as effectively and gets less out of the individual capabilities of its team members, which reduces the contribution of collective knowledge to performance. Thus, Case #3 falls short of the Base Case’s performance. Although recovering slightly after about half of the simulation period (week 112), its recovery is weak and the performance level remains below the Base Case throughout the simulation.




Figure 6. Case #3 compared to Base Case.

Notes. To create the two traces shown in Figure 6, we varied two constants. The individual learning rate λ_i is 1.0 unit of knowledge per completed task in the Base Case and 2.0 in the Case #3 simulation. The collective forgetting rate χ^c is 1/36 or 2.78% per week in the Base Case and double (5.55% per week) in the Case #3 simulation. During the 240-week simulation, Case #3 has the following summary statistics: Average = 2.6; Maximum = 8.0; Final value = 1.5 [units: tasks completed per week].

Case #4. Last, we consider a case in which AI increases both the collective forgetting rate χ^c and the individual forgetting rate χ^i . It represents a case in which AI disrupts the group's role structure and causes obsolescence of individual knowledge.

For instance, the use of surgical robots to assist in surgery can alter the practices of operating teams in a way that renders prior learning strategies of surgical trainees ineffective, since trainees no longer assist the surgeons to the same extent they did before in robot-assisted surgery (Beane, 2019). A study in the banking industry showed that the skills by junior investment bankers are rendered obsolete in areas where the analyses that junior bankers previously completed are now automated by AI tools (Anthony, 2021). When AI autonomously completes part of the overall analysis, not only is individual knowledge depreciated, AI also changes the interaction patterns between technology, junior-, and senior bankers (Anthony, 2021). Teams faced with robotic innovations (a pharmacy dispensing robot) also need to re-negotiate disrupted role structures when the relative skills or status between occupational groups changes (Barrett et al., 2012). Teams need to rebuild collective and individual knowledge continuously when they are faced with a technology that renders both individual and collective knowledge obsolete.

When AI disrupts an existing role structure, knowledge of “who knows what” in the team—may unwittingly be disrupted (seen in Cases #3 and #4). We interviewed a senior software engineer at a top 3 AI firm (according to market capitalization) who works on multiple applications and is “tech lead” for one of them. Software engineers at this firm each work individually on component “services” that are modules of code. Those component services need to be integrated at a higher level to create an application. It is crucial to align individuals to create component services that will work together coherently and robustly at the application level. If an individual programmer does their “own thing,” their services will not be aligned with the overall application's goals. To succeed, management and ideally other programmers need to “track what the others are doing.” AI has

not only enabled individuals' activities "to change faster, but it also [the engineer threw up her hands up at this point] mutates unpredictably." The result is that while AI has accelerated individuals' productivity, the hoped-for overall performance increases in creating applications have not yet happened.

Asked how AI might help overcome problems of "mutation" in the role structure, the senior software engineer commented that "it would help us deal with the team's treadmill of rework keeping us from getting the apps to work." In particular, she commented that if an AI could help the team members document each individual's work on services and then summarize that documentation so that managers and other individuals could more easily understand others' roles, that "would be golden." This suggests that AI with some agency over managing, tracking, and analyzing the role structure could help avoid these integration problems (see Case #1).

In another interview with a software developer of a major software company using AI coding assistants, the software developer was less optimistic regarding the possibility that AI could take on tasks previously performed by human coders. He said that "AI isn't achieving the productivity savings we want, and we need them because we don't have enough 'hands' [qualified programmers] so that we can scale [keep up with the exponentially growing need to provide more AI capabilities] or pivot [adapt to changing market conditions]." The previous senior software engineer echoed this theme in her interview. While she did believe that AI was reducing the rate at which they needed to add headcount, her team was still being overwhelmed by scaling, which is why she was so interested in a team-level AI that could manage roles. However, she added an interesting nuance:

If a [team-level] AI actually worked out and could reduce headcount, it'll start separating out programmers . . . If anything, we need the really, really capable programmers to do more high-level work like trying to figure out what we need to do to scale over the next five years.

In short, AI might require more- rather than less-trained workers. This observation addresses a comparison of Case #3 and #4. Even though Case #4 resulted in higher long-term performance, certain organizational imperatives like rapid scaling may make replacing members (or human knowledge and skills in general) untenable.

Figure 7 compares Case #4 with the Base Case. In the short run, Case #4 is similar to the previously discussed Case #3. Because AI increases the collective forgetting rate and slows down the accumulation of collective knowledge, Case #4 initially also falls short of the Base Case's performance. In the longer run, however, Case #4 performs markedly better than Case #3. The increased individual learning rate in Case #3 leads to lower performance than the increased individual forgetting rate in Case #4. All other things equal, faster forgetting leads to a lowered level of individual knowledge that indirectly benefits performance, since overspecialization is alleviated, allowing a higher level of collective knowledge, whose contribution to performance outweighs the decreased contribution of individual knowledge. Therefore, performance in Case #4 reverses course and begins to climb at week 70. Interestingly, the indirect benefit of individual forgetting stabilizes performance even over the Base Case after about week 120.

Case comparison. Figure 8 shows all cases—that have been shown separately (in Figures 4 to 7)—in one graph. At the beginning of the simulation, up until about week 75, the visual traces of the four cases are split into two groups, according to whether AI increases the collective learning rate (Cases #1 and #2) or AI increases the collective forgetting rate (Cases #3 and #4), that is, "by row" in Table 2. The two cases where collective learning is accelerated (Case #1 and #2) lie above the Base Case; the two cases where collective forgetting is accelerated (Cases #3 and #4) lie




Figure 7. Case #4 compared to Base Case.

Notes. To create the two traces shown in Figure 7, we varied two constants. The individual forgetting rate χ^i is $1/24$ or 4.17% per week in the Base Case and double (8.33% per week) in the Case #4 simulation. The collective forgetting rate χ^c is $1/36$ or 2.78% per week in the Base Case and double (5.55% per week) in the Case #4 simulation. During the 240-week simulation, Case #4 has the following summary statistics: Average = 3.8; Maximum = 8.9; Final value = 3.0 [units: tasks completed per week].

below the Base Case. This points to the importance of collective learning: When the accumulation of collective knowledge is disturbed, the learning system's performance suffers.

All cases show the characteristic “overshoot” behavior, that is, a decline in performance after the initial peak. However, while Cases #2 and #4 stabilize on a higher level (above the Base Case), Cases #1 and #3 oscillate on a lower level (below the Base Case). This observation highlights the importance of the increased individual *forgetting* rate in the long term. In AI-Human Learning Systems where AI also increases individual learning indirectly—by increasing the number of completed tasks from which humans learn—the increased forgetting rate remedies some of the detrimental effects of overspecialization, leading to higher overall performance in the long term.

Overview of sensitivity testing. While our results for the Base Case and the computational experiments were simulated with the parameters shown in Table 1, the simulation results qualitatively hold under a wide range of parameter values. This is discussed in more detail in the Monte Carlo analyses presented in Appendix 2. For example, we ran 1000 simulations for each of the analyses presented in Figures 4 to 7. Each simulation run draws a random value for the parameters that were varied in the computational experiments (individual and collective learning and forgetting rates) from a range of values that starts 50% lower and ends at a 50% higher increase than the increase in the main analyses. For instance, when AI increases individual learning (in Case #1 and Case #3), we use an increase from 1.0 to 2.0 (units of knowledge accumulated per task completed) in the main analyses; in the Monte Carlo simulations, this parameter is randomly selected for each sensitivity run from a range between 1.5 and 2.5. We visualize the results in Figure 10, which shows that the behaviors' shapes and changes in performance versus the Base Case remain qualitatively the same in all cases.




Figure 8. Comparison of all four cases to each other and to Base Case.

Discussion

Computer scientist and psychologist Licklider (1960: 4), a pioneer of the Internet architecture, interactive computing, and AI, expressed the hope that “human brains and computing machines will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today.”

However, despite recent interest in AI “augmentation” (cf. Baer et al., 2024) and “teaming” with AI (cf. O’Neill et al., 2023) scant research is devoted to how AI use in organizations affects the human systems in which AI is embedded. In considering AI as a learning entity that is embedded in a system of work performed by humans, we advance the field’s understanding of a *systems view* of AI. While some papers have considered AI and its relation to individual workers (e.g. Wang et al., 2023), none has considered how AI affects and is affected by the *collective* learning and memory that is so vital to human learning, memory, and coordination. Organizational outcomes are often produced by workers who must learn, interpret, and recall knowledge *together*. We analyzed how using AI in a collective setting has complex dynamic effects on individual and collective learning, and the rates at which learning entities accumulate or forget knowledge that contributes to performance. We interpret these effects to better understand how AI will affect organizational systems and how AI can be managed to provide capabilities underlying competitive advantage.

Key findings

“[O]rganizational system[s] for collective encoding, storing, and retrieving knowledge” (Argote and Ren, 2012: 1375) have been proposed to constitute a microfoundation of dynamic capabilities and, ultimately, an important source of competitive advantage (Argote and Ren, 2012). These systems develop as individuals learn through experience working together (Argote and Ren, 2012: 1378) and increasingly through experience working in systems of humans, groups, and organizations in which AI is deeply embedded.

First, our results highlight the importance of collective knowledge for coordinating members with disparate expertise in teams and larger organizational systems in which AI is integrated. To date, the role of the collective (group, unit, or organization) in producing knowledge and performance is largely absent in the literature on AI. Because collective and individual learning are intertwined, each depends on the other. Without a stock of collective knowledge, individuals' learning slows, which reduces the overall performance of the system. Second, we consider the role of "forgetting" in the AI-Human Learning System. Forgetting can stem from knowledge obsolescence or from AI taking over (parts of) a task previously performed by humans. The promise of AI to take over tasks previously performed by humans is often touted as a key productivity benefit of AI. Our study adds nuance to this notion by showing why this benefit might not be realized. We find different effects of forgetting, depending on whether the forgetting occurs at the individual level or at the collective level.

The Base Case simulation integrates AI learning with individual and collective learning. Simulations of this learning system show an appreciable improvement in performance over human-only learning systems (cf. Anderson and Lewis, 2014). While not surprising, the Base Case does align with the expectation that AI helps organizations by expanding capabilities. It also highlights the dangers of expanding individual knowledge unchecked until individuals become overspecialized, leading to problems with communication and collaboration among members. Insufficient collective knowledge has dramatic and long-lasting effects on the learning system. Hence, using an AI can improve performance overall relative to not using it, but it does not necessarily prevent performance collapses when team coordination breaks down.

Computational experiments were conducted to compare different cases, representing the potential effects of AI under four different conditions. The cases reflect examples in the literature and technology already used by organizations. Case #1 shows that teams generally benefit from AI that increases individual and collective learning but also highlights the importance of considering the indirect effects on learning. While performance is higher than in the Base Case for most of the simulation, the direct increase of individual and collective learning is countered by AI's indirect effects on the learning system and, ultimately, performance declines in the long run. With Case #2, the computational experiments revealed a somewhat counterintuitive and surprising development in the long run. An AI that increases individual forgetting instead of individual learning prevents some of the decline in performance and, in the long run, might lead to higher (above Base Case) performance.⁸ Cases 3 and 4 both analyze an AI that increases collective forgetting. Simulations show that when AI increases collective forgetting, early performance is significantly lower than in the Base Case. In the long run, however, increased performance is observed when AI increases individual forgetting (Case #4) compared with when AI increases individual learning (Case #3).

Theoretical implications

These findings substantiate recent conceptual work that describes technologies as more constitutive and questions whether humans performing tasks—who used to be alone at the center of theories of organizing—can be considered in isolation from the team structure in which they are embedded alongside AI (Faraj and Leonardi, 2022). Our results point toward intricate ways that AI as a learning entity is intertwined in relationships with human team members. While the role of the collective in producing knowledge and performance is largely absent in the literature on AI, the results from our simulations highlight its importance in coordinating members with disparate expertise. Instead, collective and individual learning are entangled in a complex dynamic dance with AI. If AI inhibits the accumulation of collective knowledge, any benefits from augmenting

individual learning are not realized. Indeed, an AI that causes the obsolescence of individual knowledge is often superior in such cases. In contrast, AI with some agency over important pillars of team functioning, such as the division of cognitive labor and role assignment (Cases 1 and 2), can improve performance. These results highlight how AI is central to organizing and organization theory.

Our findings contribute to organizational learning theory and how it relates to the microfoundations of organizational capabilities and competitive advantage (Argote and Ren, 2012; Volberda et al., 2021) providing a link between decisions of senior managers and organizational performance (Argote and Levine, 2020; Harvey et al., 2022). In support of this view, Argote (2013) noted that organizations adapt through learning activities, which occur within and between teams in a larger organizational context. Strategy scholars generally agree that managers must leverage knowledge generated by teams to gain and maintain competitive advantage (e.g. Eisenhardt and Martin, 2000), but without a much-needed focus on learning at the team or unit level (Harvey et al., 2022). Our study shows how organizational capabilities can develop and change, by linking AI to individual- and team-level learning processes.

These insights may help to better understand how organizations can establish a sustainable competitive advantage with AI when many companies have access to similar technologies (Kemp, 2024; Zeng and Glaister, 2018). The field of strategic organization has defined its focus as the “design, administration, arrangement, or structuring of [a . . .] group, or system so that it will be most useful or have the greatest effect” (Baum et al., 2022: 683). Our study emphasizes the importance of developing systems that maintain a shared understanding, promoting cross-functional communication. When AI is integrated in teams with varying expertise and the teams’ patterns of interactions, AI’s learning and actions become tailored to the specific contexts and teams. Such a socially complex learning system—once established—is difficult to imitate by competitors (Argote and Ren, 2012; Barney, 1991). The flip side is that this also makes it challenging for strategic leaders to establish and manage such a system.

Our results also inform another core area of (behavioral) strategy: the consideration of dual (or multiple) opposing goals (e.g. Levinthal and Rerup, 2021). In designing learning systems, the delivery of high performance should be considered alongside the need for humans to accumulate more expertise than may be optimal for (short- or medium-term) performance. For instance, do our results mean that human expertise is losing some of its value as it is—at least in part—replaced by AI? Not necessarily. Studies on automation technologies (e.g. Bainbridge, 1983; Baxter et al., 2012) have shown time and again that it is problematic when humans are “out of the loop” since their knowledge deteriorates, but is necessary to intervene when technology reaches its limits. Similarly, human expertise might play an important role in monitoring AI performance, especially if transparency of AI’s work products is low or new contexts make AI unreliable.

Practical implications

Strategy practitioners must consider and integrate AI when (re-)designing organizational systems. We advocate that collective knowledge should be comprehended as an organizational design principle. When considering the “intelligence” of emerging technologies, the focus usually is on their ability to solve tasks that, previously, only humans were able to solve. An (at least) equally important dimension in organizational design is the social embeddedness of technology. When integrating AI into organizations—and while considering AI’s implications for individual and collective learning—tasks and organizational structures should be designed accordingly. For example, modular interfaces could create sub-tasks where AI can be used in isolation from team collaboration.

Other parts of the task, where teams need more collective knowledge, are “protected” and AI use is restricted in ways that allows for building individual and collective knowledge. Although at the expense of (short-term) performance, this would benefit longer-term performance by setting the learning system on a more beneficial long-term path.

Another practical implication of our study is that using AI to improve individual performance in isolation, in a broad range of contexts is likely insufficient to improve overall performance. Without consideration of the collective level, AI might solely improve the individual contribution to performance, interfering with coordination between individuals, often in unpredictable ways. This makes it more difficult to track and understand how role structures change over time, inadvertently creating higher level misalignments. To leverage new configurations of people and AI, investments must also be directed toward systemic integration in contrast to merely providing “tools” for individual workers.

Given that in the cases we investigated AI generally improves peak performance in the short run, an important lever to optimize performance is finding ways to prevent the declines that follow in the longer run. Organizations need to actively monitor and strategically manage their learning system’s status and dynamically react to overspecialization or other warning signals of discoordination. Might a sudden adaptation like disbanding and reassembling teams—and the corresponding changes in collective and individual knowledge—help to reset the system back to a state with higher performance? How often and following what developments are these adjustments necessary?

Limitations and future research

While our results are robust to a wide range of parameters (see the Monte Carlo analysis in Appendix 2), our model cannot consider all variables that have been examined in AI research. For instance, we do not focus on *transparency* or *reliability* of AI and their effects on the learning system, or how the affective responses to AI (e.g. aversion, liking, identification) might influence how and what individuals learn. Future studies that examine how these characteristics affect the AI-Human Learning System would advance current research.

Another promising area for future research could consider the social attributions connected to AI use, as these could also affect learning and expertise. Studies on algorithmic management indicate that teams supervised by algorithms are viewed as less capable and that these perceptions are negatively related to the budget that is allocated to the teams (Schweitzer and De Cremer, 2024).

Other intriguing avenues for future research are the strategic considerations in the design and management of these learning systems. In our model, AI learning is affected directly by task completion and indirectly by the particular state of learning at the individual and collective levels. We do not consider changes or tweaks to the AI that might be made from within or outside (e.g. central unit in charge of AI development) of the group. Future research could consider human interventions that may alter the Human-AI Learning System. Such interventions might include task changes that reset parts of the learning system in a way that avoids the longer-term performance decline. As certain task changes require the team to build new knowledge and co-specialize in different ways, the pitfalls of overspecialization and disrupted role structures can be managed. An exciting extension of our work could also be the consideration of tasks that require other learning types and innovation activities, such as exploitative and explorative search (March, 1991).

Finally, any model is necessarily a simplification of reality and is, therefore, incomplete. For example, in order to maintain focus and readability, we have not considered the effects of the accumulations of knowledge resulting in other aspects of task completions such as improving quality (although this is often modeled similarly, see the literature review in Anderson and Chandrasekaran, 2024). Such simplifications nevertheless produce complex dynamic results that are nearly

impossible to predict in advance. That said, the computational model and simulations have revealed several insights that contribute to theory and practice and may further guide future research on AI-Human Learning Systems.

Conclusions

AI-enabled technologies are rapidly evolving. It seems likely that AI will increasingly be integrated organization-wide and become more interdependent with the work and knowledge of human knowledge workers and teams. Recent trends such as “AI agents” and “multi-agent AI systems” also point to these technologies becoming more autonomous, with more leeway to “act” alongside human workers. If true, a systems view that considers AI as embedded in organizational work is even more important. With AI becoming more autonomous, carefully designing AI-Human Learning Systems is increasingly important for leveraging human knowledge and effectively integrating AI into organizations.

ORCID iDs

Patrick Figge  <https://orcid.org/0000-0002-4418-6033>

Edward Anderson  <https://orcid.org/0000-0002-7747-1607>

Kyle Lewis  <https://orcid.org/0000-0002-6922-4876>

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Notes

1. We use the terms “teams” and “groups” interchangeably.
2. We use the terms performance and productivity interchangeably. When we refer to performance in the mathematical model, it is understood that performance is a function of task completion over time.
3. We note that artificial intelligence (AI) can also learn from limited case data (e.g. Palmarini and Mitchell, 2024) and can make inferences based on auxiliary information such as descriptions, attributes, or taxonomies (e.g. Webb et al., 2023). In the current study we focus on AI learning that occurs as a function of task “experience.”
4. For those familiar with the system dynamics methodology, we also present a stock and flow visualization of the model in Appendix 1.
5. See Argote (2024) and Anderson et al. (2019) for reviews of the learning curve formulation in the organizational theory and operations management literatures, respectively.
6. In the system dynamics literature, this equation would typically be written in differential equation form as $\frac{dK^i(t)}{dt} = \lambda_i l^i(t) q(t) - x^i K^i(t)$. However, because the values are calculated only once each time period, the equations written in the discrete K_t^i notation are computationally identical to the differential equation when using the standard Euler integration method. We use the discrete notation in this article because it is more familiar to many readers.
7. Anderson and Lewis (2014) presented a learning system model that includes individual and collective learning *without* AI (see pp. 370–371). All parameters in the “stability model” (Anderson and Lewis 2014) are identical to the model simulated here, but without AI. The peak performance in the “stability” model is about 4.0 tasks/month, occurring around month 12 (equivalent to week 48).

8. Because this was somewhat counterintuitive, we investigated this further in the Monte Carlo analyses and found that this was not a result of the Base Case's particular parameterization but that this happened in almost half of all simulations.

References

- Agrawal A, Gans J and Goldfarb A (2021) *AI Adoption and System-Wide Change*. Cambridge, MA: National Bureau of Economic Research (NBER). Available at: <http://www.nber.org/papers/w28811.pdf> (accessed 24 January 2024).
- Agrawal A, Gans J and Goldfarb A (2022) *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Boston, MA: Harvard Business Review Press.
- Allen R and Choudhury PR (2022) Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion. *Organization Science* 33(1): 149–169.
- Anderson EG and Chandrasekaran A (2024) How to sustain healthcare process improvement in heterogeneous teams? Evidence from a systems dynamic model. *Production and Operations Management* 33(4): 880–902.
- Anderson EG and Lewis K (2014) A dynamic model of individual and collective learning amid disruption. *Organization Science* 25(2): 356–376.
- Anderson EG and Parker GG (2002) The effect of learning on the make/buy decision. *Production and Operations Management* 11(3): 313–339.
- Anderson EG, Jiang X, Parker GG, et al. (2019) Systems integration and the dynamics of partial outsourcing. *Production and Operations Management* 28(2): 319–340.
- Anderson EG, Parker G and Tan-Erciyes B (in press) Generative AI coding: Fast gains, hidden costs. *MIT Sloan Management Review*.
- Anthony C (2021) When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies. *Administrative Science Quarterly* 66(4): 1173–1212.
- Anthony C, Bechky BA and Fayard A-L (2023) “Collaborating” with AI: Taking a system view to explore the future of work. *Organization Science* 34(5): 1672–1694.
- Argote L (1993) Group and organizational learning curves: Individual, system and environmental components. *British Journal of Social Psychology* 32(1): 31–51.
- Argote L (2013) *Organizational Learning: Creating, Retaining and Transferring Knowledge*, 2nd edn. Boston, MA: Springer.
- Argote L (2024) Group dynamics meets organizational learning: Reflections on research. *Small Group Research* 56: 524–552.
- Argote L and Epple D (1990) Learning curves in manufacturing. *Science* 247(4945): 920–924.
- Argote L and Levine JM (eds) (2020) *The Oxford Handbook of Group and Organizational Learning*, 1st edn. Oxford: Oxford University Press.
- Argote L and Miron-Spektor E (2011) Organizational learning: From experience to knowledge. *Organization Science* 22(5): 1123–1137.
- Argote L and Ren Y (2012) Transactive memory systems: A microfoundation of dynamic capabilities. *Journal of Management Studies* 49(8): 1375–1382.
- Argote L, Lee S and Park J (2021) Organizational learning processes and outcomes: Major findings and future research directions. *Management Science* 67(9): 5399–5429.
- Bachrach DG, Lewis K, Kim Y, et al. (2019) Transactive memory systems in context: A meta-analytic examination of contextual factors in transactive memory systems development and team performance. *The Journal of Applied Psychology* 104(3): 464–493.
- Baer I, Waardenburg L and Huysman M (2024) What is augmented? A metanarrative review of AI-based augmentation. *Journal of the Association for Information Systems* 26: 760–798.
- Bailey DE, Faraj S, Hinds PJ, et al. (2022) We are all theorists of technology now: A relational perspective on emerging technology and organizing. *Organization Science* 33(1): 1–18.
- Bainbridge L (1983) Ironies of automation. *Automatica* 19(6): 775–779.
- Barney J (1991) Firm resources and sustained competitive advantage. *Journal of Management* 17(1): 99–120.

- Barrett M, Oborn E, Orlikowski WJ, et al. (2012) Reconfiguring boundary relations: Robotic innovations in pharmacy work. *Organization Science* 23(5): 1448–1466.
- Baum JA, Greenwood R and Jennings PD (2022) Constructing strategic organization: A field whose time has come. *Strategic Organization* 20(4): 683–697.
- Baxter G, Rooksby J, Wang Y, et al. (2012) The ironies of automation: Still going strong at 30? In: *Proceedings of the 30th European conference on cognitive ergonomics*, Edinburgh, 28–31 August, pp. 65–71. ACM. Available at: <https://dl.acm.org/doi/10.1145/2448136.2448149> (accessed 18 March 2023).
- Beane M (2019) Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly* 64(1): 87–123.
- Beane M and Orlikowski WJ (2015) What difference does a robot make? The material enactment of distributed coordination. *Organization Science* 26(6): 1553–1573.
- Boone T, Ganeshan R and Hicks RL (2008) Learning and knowledge depreciation in professional services. *Management Science* 54(7): 1231–1236.
- Clark HH (1985) Language use and language users. In: Lindzey G and Aronson E (eds) *Handbook of Social Psychology*. New York: Random House, pp. 179–231.
- Clark HH (1998) Communal lexicons. In: Malmkjær K and Williams J (eds) *Context in Language Learning and Language Understanding*. Cambridge: Cambridge University Press, pp. 64–87.
- Crossan M, Vera D and Pathak S (2021) Organizational learning. In: Duhaime IM, Hitt MA and Lyles MA (eds) *Strategic Management: State of the Field and Its Future*. New York: Oxford University Press, pp. 467–484.
- Cui ZK, Demirel M, Jaffe S, et al. (2024) *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*. Available at: <https://papers.ssrn.com/abstract=4945566> (accessed 5 July 2025).
- Darr ED, Argote L and Eppl D (1995) The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science* 41(11): 1750–1762.
- Davenport TH (2018) *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*. Cambridge, MA; London: The MIT Press.
- Dell'Acqua F, McFowland E, Mollick ER, et al. (2023) *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*. Available at: <https://www.ssrn.com/abstract=4573321> (accessed 5 July 2025).
- Duggan J, Sherman U, Carbery R, et al. (2020) Algorithmic management and app-work in the gig economy: A research agenda for employment relations and HRM. *Human Resource Management Journal* 30(1): 114–132.
- Edmondson AC (2002) The local and variegated nature of learning in organizations: A group-level perspective. *Organization Science* 13(2): 128–146.
- Eisenhardt KM and Martin JA (2000) Dynamic capabilities: What are they? *Strategic Management Journal* 21(10–11): 1105–1121.
- Ellis HC (1965) *The Transfer of Learning*. New York: Macmillan.
- Faraj S and Leonardi PM (2022) Strategic organization in the digital age: Rethinking the concept of technology. *Strategic Organization* 20(4): 771–785.
- Fiore SM and Wiltshire TJ (2016) Technology as teammate: Examining the role of external cognition in support of team cognitive processes. *Frontiers in Psychology* 7: 1531.
- Fraidin SN (2004) When is one head better than two? Interdependent information in group decision making. *Organizational Behavior and Human Decision Processes* 93(2): 102–113.
- Gupta P, Kim YJ, Glikson E, et al. (2024) Using digital nudges to enhance collective intelligence in online collaboration: Insights from unexpected outcomes. *MIS Quarterly* 48(1): 393–408.
- Harvey J-F, Bresman H, Edmondson AC, et al. (2022) A strategic view of team learning in organizations. *Academy of Management Annals* 16(2): 476–507.
- Hollingshead AB (2001) Cognitive interdependence and convergent expectations in transactive memory. *Journal of Personality and Social Psychology* 81(6): 1080–1089.
- Hu Q, Lu Y, Pan Z, et al. (2021) Can AI artifacts influence human cognition? The effects of artificial autonomy in intelligent personal assistants. *International Journal of Information Management* 56: 102250.

- Huber GP (1991) Organizational learning: The contributing processes and the literatures. *Organization Science* 2(1): 88–115.
- Jung MF and Hinds P (2018) Robots in the wild: A time for more robust theories of human-robot interaction. *ACM Transactions on Human-Robot Interaction* 7(1): 2.
- Kellogg KC, Valentine MA and Christin A (2020) Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14(1): 366–410.
- Kemp A (2024) Competitive advantage through artificial intelligence: Toward a theory of situated AI. *Academy of Management Review* 49(3): 618–635.
- Kozlowski SWJ and Klein KJ (2000) A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes. In: Klein KJ and Kozlowski SWJ (eds) *Multilevel Theory, Research, and Methods in Organizations: Foundations, Extensions, and New Directions*. San Francisco: Jossey-Bass, pp. 3–90.
- Krauss RM and Fussell SR (1996) Social psychological models of interpersonal communication. In: Higgins ET and Kruglanski AW (eds) *Social Psychology: Handbook of Basic Principles*. New York: Guilford Press, pp. 655–701.
- Lee J-Y, Bachrach DG and Lewis K (2014) Social network ties, transactive memory, and performance in groups. *Organization Science* 25(3): 951–967.
- Levinthal DA and Rerup C (2021) The plural of goal: Learning in a world of ambiguity. *Organization Science* 32(3): 527–543.
- Lewis K (2003) Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology* 88(4): 587–604.
- Lewis K (2004) Knowledge and performance in knowledge-worker teams: A longitudinal study of transactive memory systems. *Management Science* 50(11): 1519–1533.
- Lewis K and Herndon B (2011) Transactive memory systems: Current issues and future research directions. *Organization Science* 22(5): 1254–1265.
- Lewis K, Belliveau M, Herndon B, et al. (2007) Group cognition, membership change, and performance: Investigating the benefits and detriments of collective knowledge. *Organizational Behavior and Human Decision Processes* 103(2): 159–178.
- Lewis K, Lange D and Gillis L (2005) Transactive memory systems, learning, and learning transfer. *Organization Science* 16(6): 581–598.
- Licklider JCR (1960) Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics* HFE-1(1): 4–11.
- McNeese NJ, Flathmann C, O'Neill TA, et al. (2023) Stepping out of the shadow of human-human teaming: Crafting a unique identity for human-autonomy teams. *Computers in Human Behavior* 148: 107874.
- Malone TW (2018) *Superminds: The Surprising Power of People and Computers Thinking Together*. Boston, MA: Little, Brown and Company.
- March JG (1991) Exploration and exploitation in organizational learning. *Organization Science* 2(1): 71–87.
- Meadows D (2008) *Thinking in Systems* (Illustrated edition). White River Junction, VT: Chelsea Green Publishing Co.
- Nembhard DA and Uzumeri MV (2000) An individual-based description of learning within an organization. *IEEE Transactions on Engineering Management* 47(3): 370–378.
- Noy S and Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654): 187–192.
- O'Neill TA, Flathmann C, McNeese NJ, et al. (2023) Human-autonomy teaming: Need for a guiding team-based framework? *Computers in Human Behavior* 146: 107762.
- O'Neill TA, McNeese NJ, Barron A, et al. (2022) Human-autonomy teaming: A review and analysis of the empirical literature. *Human Factors* 64(5): 904–938.
- Palmarini AB and Mitchell M (2024) Abstract understanding of core-knowledge concepts: Humans vs. LLMs. In: *Proceedings of the 41st international conference on machine learning (PMLR)*, Vienna, 21–27 July.
- Parent-Rochelleau X and Parker SK (2022) Algorithms as work designers: How algorithmic management influences the design of jobs. *Human Resource Management Review* 32(3): 100838.

- Rahmandad H (2012) Impact of growth opportunities and competition on firm-level capability development trade-offs. *Organization Science* 23(1): 138–154.
- Raisch S and Fomina K (2024) Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review*. Epub ahead of print 5 January. DOI: 10.5465/amr.2021.0421.
- Reagans R, Argote L and Brooks D (2005) Individual experience and experience working together: Predicting learning rates from knowing who knows what and knowing how to work together. *Management Science* 51(6): 869–881.
- Ren Y and Argote L (2011) Transactive memory systems 1985–2010: An integrative framework of key dimensions, antecedents, and consequences. *Academy of Management Annals* 5(1): 189–229.
- Repenning NP (2002) A simulation-based approach to understanding the dynamics of innovation implementation. *Organization Science* 13(2): 109–127.
- Repenning NP and Serman JD (2001) Nobody ever gets credit for fixing problems that never happened: Creating and sustaining process improvement. *California Management Review* 43(4): 64–88.
- Schweitzer S and De Cremer D (2024) When being managed by technology: Does algorithmic management affect perceptions of workers' creative capacities? *Academy of Management Discoveries* 10(3): 375–392.
- Sebo S, Stoll B, Scassellati B, et al. (2020) Robots in groups and teams: A literature review. *Proceedings of the ACM on Human-Computer Interaction* 4: 176.
- Sergeeva AV, Faraj S and Huysman M (2020) Losing touch: An embodiment perspective on coordination in robotic surgery. *Organization Science* 31(5): 1248–1271.
- Shapira Z (2020) Psychological foundations of group and organizational learning. In: Argote L and Levine JM (eds) *The Oxford Handbook of Group and Organizational Learning*. Oxford: Oxford University Press, pp. 21–34.
- Serman JD (2001) System dynamics modeling: Tools for learning in a complex world. *California Management Review* 43(4): 8–25.
- Serman JD (2010) *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Columbus, OH: Tata McGraw-Hill Edition.
- Serman JD, Repenning NP and Kofman F (1997) Unanticipated side effects of successful quality programs: Exploring a paradox of organizational improvement. *Management Science* 43(4): 503–521.
- Sturm T, Gerlach J, Pumplun L, et al. (2021) Coordinating human and machine learning for effective organization learning. *MIS Quarterly* 45(3): 1581–1602.
- Traeger ML, Strohkorb Sebo S, Jung M, et al. (2020) Vulnerable robots positively shape human conversational dynamics in a human–Robot team. *Proceedings of the National Academy of Sciences of the United States of America* 117(12): 6370–6375.
- Volberda HW, Schneidmuller T and Zadeh T (2021) Knowledge and innovation: From path dependency toward managerial agency. In: Duhaime IM, Hitt MA and Lyles MA (eds) *Strategic Management: State of the Field and Its Future*. New York: Oxford University Press, pp. 445–466.
- Wang W, Gao G and Agarwal R (2023) Friend or foe? Teaming between artificial intelligence and workers with variation in experience. *Management Science* 70: 5753–5775.
- Webb T, Holyoak KJ and Lu H (2023) Emergent analogical reasoning in large language models. *Nature Human Behaviour* 7(9): 1526–1541.
- Wegner DM, Giuliano T and Hertel PT (1985) Cognitive interdependence in close relationships. In: Ickes W (ed.) *Compatible and Incompatible Relationships*. New York: Springer, pp. 253–276.
- Wilson HJ and Daugherty PR (2018) Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review*, July–August. Available at: <https://hbsp.harvard.edu/product/R1804J-PDF-ENG>
- Wittenbaum GM, Stasser G and Merry CJ (1996) Tacit coordination in anticipation of small group task completion. *Journal of Experimental Social Psychology* 32(2): 129–152.
- Wright TP (1936) Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences* 3(4): 122–128.
- You S and Robert L (2017) Teaming up with robots: An IMOI (inputs-mediators-outputs-inputs) framework of human-robot teamwork. *International Journal of Robotic Engineering* 2: 1–7.
- Zeng J and Glaister KW (2018) Value creation from big data: Looking inside the black box. *Strategic Organization* 16(2): 105–140.

Author biographies

Patrick Figge is a postdoctoral researcher and lecturer at the University of Passau. His research focuses on innovation, entrepreneurship, team science, and the strategic management of technologies.

Edward Anderson is the Betty and Glenn Mortimer Centennial Professor for Business at the University of Texas McCombs School of Business and the University of Texas Supply Chain Management Center. His research focuses on supply chain innovation, the economics of platform and information system design, and distributed product development.

Kyle Lewis is Professor Emerita and was the Christian A. Felipe Chair in Technology Management at the University of California, Santa Barbara. Her research examines how organizations leverage individual and collective knowledge, with a focus on the performance of teams engaged in knowledge work.

Appendix I

Stock and flow diagram

Figure 9 presents a stock-and-flow diagram of the model. Note that there are some “dummy” constants in the diagram, the purpose of which is purely to avoid unit check issues. They are all set to a value of 1.0. Also, there are three “forgetting rate multipliers” that are all set to 1.0 and varied only during sensitivity testing.




Figure 9. Stock-and-flow diagram.

Appendix 2

Sensitivity testing with Monte Carlo analyses

In addition to the standard robustness tests generally performed on system dynamics-based models (Sterman, 2010), we also conducted numerical analyses over and above the ones presented in the manuscript to check the robustness of our findings because the learning curve and forgetting parameters can be expected to vary widely from organization to organization and even from (task) context to context.

First, we conducted Monte Carlo analyses using uniform distributions that vary the key parameters studied in the main analyses to conduct the computational experiments that investigate artificial intelligence's (AI) influence on the AI-Human Learning System that we conceptualized. For each of these we ran the simulations with different random combinations of parameter values between minimum and maximum sensitivity values.

In the main text, the results from the main analyses are visualized in Figures 4 to 7. These four cases represent the theoretically conceptualized ways in which an AI can affect human learning in an AI-Human Learning System. We conducted computational experiments to compare these four cases with a Base Case simulation. Table 2 provides an overview of these four cases that were generated by considering AI's influence on all inflows and outflows into the human knowledge stocks: the individual learning rate λ_i , the individual forgetting rate χ^i , the collective learning rate λ_c , and the collective forgetting rate χ^c . In the main analyses, we increase the learning rates from 1.0 in the Base Case to 2.0 units of knowledge accumulated per task completed in the computational simulations (in the cases where learning rates are affected). For the Monte Carlo sensitivity analyses, we consider a minimum sensitivity value of 1.5 and a maximum sensitivity value of 2.5, that is, we investigate how our results change when AI increases the learning rate by 50% less or 50% more than in the main analyses. We set the minimum and maximum sensitivity value for the forgetting rates accordingly: We multiplied the forgetting rates by a factor drawn from the interval of [1.5, 2.5], instead of the doubling of the forgetting rate that was used to model an increase in the forgetting rate in the main analyses (in the cases where forgetting rates are affected).

Figure 10 shows the corresponding graphs depicting the sensitivity analyses for the four cases. Each of the four graphs is based on 1000 simulation runs with different random combinations for the individual and collective learning and forgetting rates from a uniform distribution between minimum and maximum sensitivity values just described. As in the main analyses of the four cases, all other parameters remain at the base values as listed in Table 1 in the main text.

In each graph, 75% of the simulation runs fall within the green shaded area. When considering the expanded area shaded in blue, this increases to 90%. The median value of all the simulations at each week in the simulation is shown by the black line.

When comparing these four graphs in Figure 10 with Figures 4 to 7 in the main text, we see that the median line is almost identical to the single simulation run representing each of the four cases. In addition, the green and blue shaded areas are very narrow. These areas represent the variation of these runs due to different learning and forgetting rates (i.e. "confidence intervals" around the cases presented in the main text). The confidence interval bands around the median case show that—within the broad range of values used to generate these runs—the behavior of the findings from the computational experiments is remarkably robust. The behaviors' shapes and changes in performance versus the Base Case remain qualitatively the same in all cases.

We conducted a further—and very extensive—Monte Carlo analysis that varies not only the four learning and forgetting rates used to generate the computational experiments depicted in Table 2, but also an extensive list of parameters using the sensitivity ranges shown in Table 4 below.




Figure 10. Sensitivity runs for Figures 4 to 7.

Notes. The black line (within the green area) shows the median simulation. The green area includes 75% of the 1000 simulation runs; the blue area includes 90% of the 1000 simulation runs.

Figure 11 shows the results from these analyses. As one would expect, the results vary significantly. However, the qualitative behavior including a short-term peak in performance followed by a long-term decline remains consistent for 90% of the runs indicating that it is characteristic across a very wide range of parameterizations. It is worth mentioning that the gap between the 50th and 90th percentile at peak magnitude is much wider than that between the 10th and 50th percentile and is driven by the large non-linearities in the system (especially, the learning curve) interacting with reinforcing loops. For example, as is well known from the learning curve literature (Argote, 2013), the gap in performance over time between a learning curve that increases productivity by 10% per doubling of accumulated knowledge and one that increases by 20% will be much smaller than the gap between 20% and 30%. The reason that there is a reconvergence toward the end is that the negative loop of overspecialization becomes active and begins to counter the learning curve's non-linear growth later in the simulation. One could, therefore, argue that the Base Case shown in the main text (Figure 2) is a rather conservative choice and that many organizations in practice might see more extreme versions, characterized by the same behaviors that we uncover, but with an even larger magnitude.

Next, we quantitatively investigate key patterns that characterize the results that we discuss in the main analyses. First, we assess how often an “overshoot” behavior occurs, that is, in how many of the sensitivity runs we observe that productivity reaches a short-term peak followed by a long-term decline. Then, we assess how often the peak performance of the four cases relates to the Base Case in the same way as it does in the main analyses, that is, if Case #1 and Case #2 are above and Case #3 and Case #4 below the Base Case's peak performance, as well as the extent of the




Figure 11. Sensitivity runs for Figure 2.
Notes. The black line (within the green area) shows the median simulation. The green area includes 75% of the 1000 simulation runs; the blue area includes 90% of the 1000 simulation runs. The range of the y-axis has been expanded for this graph, that is, ranges from 0 to 60 rather than 0 to 20 as the other graphs.

Table 3. Percentage of overshoot behavior from 1000 simulations.

Performance peak followed by a decline	Base Case	Case 1	Case 2	Case 3	Case 4
Yes	98.5%	98.5%	98.8%	99.7%	99.6%
No	1.5%	1.5%	1.2%	0.3%	0.4%

difference between the respective peak performances. For these analyses, we used the parameter ranges displayed in Table 5. That is, for each of 1000 simulation runs, a random value from a uniform distribution between minimum and maximum sensitivity value was drawn for each of the parameters listed in Table 5. (Note: The ranges between maximum and minimum are even broader than the ones used in the prior Monte Carlo analyses. The reason is that while the values displayed in Table 4 account for most variation that we see in most contexts in the real world, we wanted to include more extreme contexts as well.)

First, the “overshoot” behavior. In the main text, we discussed that all cases show a peak in performance early in the simulation followed by a long decline, perhaps with some recovery and oscillation after the decline. Table 3 shows how many percent of the 1000 simulations we did also find this pattern. The overwhelming majority of simulations in our sensitivity analyses replicate this pattern. Only a few outliers—between 0.3% and 1.5% of the simulation runs—are not characterized by this behavior.

Next, we discuss Figures 12 to 15, which show the ratio of the peak performance of each of the four cases divided by the peak performance of the Base Case. In the main analyses for Case #1 and Case #2, we observe a peak performance that is higher than in the Base Case. Figure 12 shows the ratio of the peak performance in Case #1 divided by the peak performance in the Base Case for 1000 sensitivity simulation runs. Both Case #1 and the Base Case are simulated using random draws from a uniform distribution along the interval shown in Table 5. The y-axis




Figure 12. Ratio of peak performance in Case #1 divided by peak performance in Base Case.




Figure 13. Ratio of peak performance in Case #2 divided by peak performance in Base Case.

shows a frequency count, that is, how many out of the 1000 simulation runs fall into a certain peak performance ratio range as indicated on the x-axis. The x-axis shows these ranges for the ratio bins. A ratio of less than 1 indicates that the Base Case peak performance is higher than Case #1’s *disconfirming* our findings in the main analyses. A ratio greater than 1 indicates that we observe a higher peak performance in Case #1 than in the Base Case, *confirming* the findings in the main analyses. In all simulation runs, we indeed observe a ratio higher than one, which indicates that our finding is robust even in extreme scenarios.




Figure 14. Ratio of peak performance in Case #3 divided by peak performance in Base Case.




Figure 15. Ratio of peak performance in Case #4 divided by peak performance in Base Case.

Figure 13 shows the corresponding frequency distribution across the 1000 sensitivity runs conducted for Case #2. Here, we see that in a small fraction of the runs (approximately 5%) the peak performance is lower than in the Base Case. For the vast majority of the runs, however, the sensitivity analyses confirm what we see in the main analyses: Case #2 has a higher peak performance than the Base Case.

Table 4. Parameter ranges for Monte Carlo analyses for Figure 11.

Parameter	Meaning	Base	Min	Max
χ^i	Individual fractional forgetting rate	1/24	$0.5 \times 1/24$	$1.5 \times 1/24$
χ^c	Collective fractional forgetting rate	1/36	$0.5 \times 1/36$	$1.5 \times 1/36$
χ^a	AI fractional forgetting rate	1/24	$0.5 \times 1/24$	$1.5 \times 1/24$
γ_i	Individual learning curve parameter	0.2	0.1	0.3
γ_c	Collective learning curve parameter	0.4	0.3	0.5
γ_a	AI learning curve parameter	0.2	0.1	0.3
λ_i	Base individual learning rate per task completed	1.0	0.5	1.5
λ_c	Base collective learning rate per task completed	1.0	0.5	1.5
λ_a	Base AI learning rate per task completed	1.0	0.5	1.5

Table 5. Parameter ranges for Monte Carlo analyses for Table 3 and Figures 12 to 15.

Parameter	Meaning	Base	Min	Max
χ^i	Individual fractional forgetting rate	1/24	$0.33 \times 1/24$	$3 \times 1/24$
χ^c	Collective fractional forgetting rate	1/36	$0.33 \times 1/36$	$3 \times 1/36$
χ^a	AI fractional forgetting rate	1/24	$0.33 \times 1/24$	$3 \times 1/24$
γ_i	Individual learning curve parameter	0.2	0.1	0.5
γ_c	Collective learning curve parameter	0.4	0.1	0.5
γ_a	AI learning curve parameter	0.2	0.1	0.3
λ_i	Base individual learning rate per task completed	1.0	0.5	1.5
λ_c	Base collective learning rate per task completed	1.0	0.5	1.5
λ_a	Base AI learning rate per task completed	1.0	0.5	1.5

Figure 14 shows results for Case #3. In this case, the bulk of the runs (greater than 95%) has a ratio of less than 1.0, which means that the peak performance of Case #3 is less than that for the Base Case, which confirms our prior findings.

Finally, Figure 15 shows the results for Case #4, which again confirms our prior findings as all runs have ratios below 1.0.