

*Preprint – published in
Harvard Business Review, October 2024*

How to Use AI to Build Your Company's Collective Intelligence

by Christoph Riedl

SUMMARY. When done right, AI can not only improve short-term productivity of organizations but can also increase their long-term performance by expanding the space of opportunities the organization considers by supporting learning, increasing intellectual diversity, and enhancing organizational culture. That is, it can increase the organization's collective intelligence. Companies should focus on three elements of collective intelligence: collective memory, collective attention, and collective reasoning. When done right, managers embrace a way of deploying AI so that skills are enhanced, curiosity is nurtured, and communication and coordination are improved.

In the management world right now, it is common to think about AI in terms of automation vs. augmentation. Whereas automation implies machines taking over tasks previously performed by humans, augmentation refers to the cooperation between humans and machines in executing tasks. According to the conventional wisdom, *augmentation* leads to higher performance than *automation* because it alleviates downsides such as short-term thinking, loss of flexibility due to lock-in, and loss of human intuition and skill. That is, downsides that counter long-term success.

Unfortunately, this view misses a key point: Augmentation doesn't avoid automation, it simply hides it, usually in some lower-level information gathering or decision-making task. For example, when organizations use a generative AI assistant to *augment* the task of writing product descriptions for an online store, they might be *automating* the task of writing first drafts. This automation replaces human intuition, expertise, experience, and reasoning with whatever the AI system decides is appropriate, and may still lead to negative effects such as deskilling. So, the downsides of automation that the augmentation strategy ostensibly helped managers avoid, are simply pushed down to a lower level.

To ensure AI deployment advances strategic goals and supports key objectives, managers can adopt a different mindset: using AI to increase the collective intelligence of the entire organization.

Collective intelligence is the shared intelligence that emerges from collaboration, collective efforts, and competition. It reflects groups' ability to achieve consensus, solve complex problems, and adapt to changing environments. Recent research suggests that collective intelligence emerges from [three interdependent ingredients](#): collective memory, collective attention, and collective reasoning. Managers can apply this idea to target specific areas in which AI can elevate the organization's collective cognitive abilities and drive more informed decision-making in ways that are human centered and amplify human creativity.

AI can support each of these three processes. Here's how.

	Purpose	The Risk <i>How AI-based automation makes intellectual spaces smaller</i>	The Opportunity <i>How AI can make intellectual spaces larger</i>
Collective Memory	Skills and knowledge of individuals; group processes to cooperatively allocate, retrieve and update knowledge	Leads to individual skill loss and rigid structures which limit flexibility, human intuition, and information exchange	AI can enhance individual's memory, and help coordinate distributed knowledge and skills. AI can help unearthing hidden information within and across individuals such as connecting teachers and learners. AI can support formation and updating of mental models of who knows what within the organization.
Collective Attention	Ability to align focus of attention to process information in the environment; group processes to allocate and retrieve attention, and update understanding of others' focus areas and current demands	Undermines individual autonomy and task variety by micromanaging, taking over decision-making processes, monitoring and automating authority	AI can enhance individual attention, limit costs of task switching, and improve joint awareness of each other's workloads and availability. AI can synchronize attention and help groups cooperatively focus, and enhancing their autonomy and identity.
Collective Reasoning	Ability to align goals, priorities and motivation; group processes that cooperatively allocate priorities, retrieve commitment, and update understanding of others' goals and motivations	Excludes diverse views, amplifies views of technical experts, introduces algorithmic monocultures	AI can enhance individuals' reasoning abilities and amplify diverse thinking styles and backgrounds. AI can help sense changes in the environment. AI can help individuals learn what group members care about, draw inferences about their goals by explaining and elaborating differences that may otherwise be difficult to understand, and prioritize diverse group goals.

Collective Memory: Enhance individual ability to learn and act on opportunities.

Generative AI is already reshaping the jobs of knowledge workers. For example, on Wall Street, financial institutions are contemplating AI to automate tasks previously done by entry level workers, such as financial analysts. As one executive at a major bank put it, “The easy idea is you just replace juniors with an A.I. tool” (NYT). While appealing for its short-term productivity gains, such a strategy poses obvious challenges for longer-term outcomes like skill and knowledge development in the organization: If entry-level jobs are fully automated, how will workers learn and become highly trained experts? Such a shift not only risks stunting the development of institutional knowledge it also disrupts the social networks and knowledge-sharing processes that are essential to access and maintain collective knowledge within an organization.

Collective memory refers to how groups allocate, retrieve, and update collective knowledge. This process allows individuals to specialize and remember different and complementary details, so that collectively the group can recall and use more available knowledge than they could individually. For this process to be effective, groups also need to be able to dynamically update their understanding of “who knows what” as they learn about each other's competencies.

AI can help companies develop, retrieve, and update collective knowledge.

For example, NVIDIA developed a [generative AI-based chatbot assistant](#) to answer questions about complex engineering tasks such as industrial chip design. The chatbot builds on existing

large language models — so-called foundation models — but has been modified with internal documents, code, and internal communication, such as emails and instant messaging. The system can explain complex design topics, helps engineers quickly find technical documents, and answers questions about designs, tools, and internal infrastructures. It thus helps make knowledge distributed throughout the organization more widely accessible. This is particularly relevant in a multinational company like NVIDIA, where human experts are often globally distributed which can make it inconvenient to find immediate help quickly.

When used this way, AI systems support the retrieval of collective knowledge held throughout the organization. This may be particularly helpful for individuals who may otherwise lack meta knowledge about where to find relevant information in the organization. By pooling knowledge in one place, the AI and the human that use it may be able to draw connections and inferences that may otherwise be difficult given human's limits to process large quantities of data. It also improves retrieval of expertise by helping identify relevant experts in the organization.

This approach can also guide the allocation of new information and tasks to individuals based on an improved understanding of expertise and specialization. Instead of automating tasks of entry level workers, AI could help them develop specialized skills and expertise by repeatedly allocating tasks that require specific expertise to the same individual, while simultaneously giving them the ability to retrieve relevant knowledge to complete those tasks.

Finally, it supports the development and updating of metaknowledge of who knows what. By drawing on communication between individuals, such systems could keep updated metamemory of changing expertise and increased specialization across the organization. As a result, individuals can more efficiently retrieve knowledge available in the organization, can channel incoming information to the appropriate person, and can keep current on who knows what. Such AI applications would support collective memory of the organization, thereby vastly expanding its collective cognitive capabilities.

Collective Attention: Shaping how individuals and groups process information.

One of the key areas in which generative AI already sees heavy use is to streamline communication. While the potential productivity gains of AI supported communication are appealing, such uses can also affect how groups coordinate and align their focus of attention.

Collective attention refers to how groups allocate and align their focus on key tasks and priorities. Sometimes this means that multiple individuals will attend to the same event (e.g., when multiple people listen as one of them speaks), while at other times it means that individuals attend to different events so as to maximize the group's collective attention. Again, groups also need to update their understanding of each other's current focus of attention and current demands.

AI can help guide the allocation of attention by streamlining communication, optimizing schedules, workflows, and to-do lists. This can enhance collective attention by limiting costs of

task switching and reduce coordination costs associated with interdependent tasks. That would free up a team's resources to focus their attention on creative matters. AI can support the retrieval of attention by supporting group routines that encourage synchronous coordination. For example, AI could help coordinate each individual members' focus so that they can attend to the [same target together](#) or divide their attention across different targets, depending on the requirements of the situation. AI can also support updating of collective attention by helping individuals develop joint awareness of each other's workload and availability. Used in this way, AI can help organizations maximize the use of their collective attentional resources.

Collective Reasoning: Help groups align their goals and priorities.

Collective reasoning refers to a group's ability to reason about each other's individual and collective goals in order to align on collective goals and priorities to maximize joint rewards. Groups need to allocate priorities and retrieve commitment of their members. Groups also need to update their understanding of each other's goals, motivations, and priorities.

AI can improve goal alignment by integrating information from different sources (e.g., using AI driven data integration). It can also enhance organizational culture by improving group processes that cooperatively allocate priorities and retrieve commitment.

To collaborate effectively, groups need to understand where each are coming from and bridge communication gaps among people. AI can help improve the group's capacity to align members around shared goals and priorities by providing summarization and analysis of individual viewpoints. Current gen AI tools have also proven useful in fostering goal alignment by making the underlying reasoning behind goals and priorities more accessible for those with different backgrounds. Generative AI in areas such as software development also offers the potential to increase collective intelligence by involving more people who otherwise couldn't be involved in the process because they lack technical background. For example, [HubSpot](#) is exploring how AI can help automate A/B tests by simplifying the process of generating website variants that the A/B test compares. Such uses of AI can enhance individuals' reasoning abilities and amplify diverse thinking styles and backgrounds.

AI can also help individuals update and learn about each other's priorities. Employees are already putting those ideas into practice with public AI chatbots as tools that aid collective reflection. [Popular anecdotes](#) of AI [use](#) at work as a tool for reflection on how to communicate with colleagues in difficult social situations. By tapping into social cues and role-play difficult social situations in this way, individuals in a group may gain access to higher-order reasoning capacities. When used in this way, AI can enhance how groups develop shared values and priorities and help regulate resources and priorities for efficient execution. Such applications of AI systems could help individuals learn what group members care about, draw inferences about their goals by explaining and elaborating differences that may otherwise be difficult to understand, and prioritize group goals. This allows them to better allocate their joint priorities by negotiating their collective goals and ensure the group pursues the most rewarding goal.

How to make your AI strategy a success

When AI is deployed as a production tool to automate processes it often acts in a way that is inherently deskilling, lead to rigid structures, and homogenize solutions, all of which limit companies' capacity to adapt and adjust to changing environments. As a result, they have the tendency to reduce the diversity of ideas, hypotheses, and processes organizations consider.

Managers can counteract this trend with a few strategies:

Use AI as a coach and coordination tool. A gen AI chatbot can automate a writing task, but it can also be used as a coach who gives feedback and critiques. In the knowledge management system developed by NVIDIA, AI acts as a support system designed to enhance the collective intelligence of the organization. Through this system, employees gain access to valuable knowledge, find opportunities to gain new skills, and update their understanding of “who knows what.” To see whether AI support systems are effective, managers can evaluate them using usage statistics, satisfaction surveys, or performance evaluations of employees.

Use AI as a production tool to supercharge experimentation and expand intellectual spaces. AI can help design experiments, identify promising drug candidates, facilitate what-if analyses, or create multiple first drafts from which the best one is selected for further development. In all these examples, AI expands the intellectual space by suggesting new and different ideas that may otherwise have been left unexplored. However, managers need to be mindful of one potential trap. When generative AI is used as a production tool that expands opportunities — such as multiple first drafts being considered — managers need to resist the urge to consider this wasteful redundancy that could be streamlined to harness short-term productivity gains.

Whether AI expands options or decreases options it depends on how it is embedded into organizational processes. When AI driven productivity gains can be leveraged to unburdening employees, increase their autonomy, and free up time for creative thinking, AI has the potential to supercharge collective intelligence. [Research on creativity](#) has long ago established the important role of access to necessary resources, such as time, money, and equipment, can facilitate the creative process. Empowering employees to adapt their roles can be a powerful pathway to ensure the intellectual space is expanded. Soliciting input from the individuals who are most directly affected and setting up systematic measurement and monitoring to bring short-term actions in line with long-term strategic goals, can help managers get this right.

Beware of AI risks that reduce cognitive diversity and amplify inequalities. Inevitably, there are also risks to adding AI into your organization's processes. For example, [research conducted in my lab](#) shows that AI can significantly affect what teams pay attention to, irrespective of the quality of the AI's contribution. When teams are joined by an AI voice assistant, they started to align their attention with that of the AI assistant, they even adopted the specific terminology used by the AI, which further shaped where groups directed their attention. What surprised us the most is that teams also adopted the language used by the AI in areas that were unrelated to the task at hand. This happened even when the AI provided only

unhelpful input and despite the fact that teams realized that the AI was unhelpful. They even reported that they didn't trust the AI. Another [study](#) we recently conducted demonstrated that while AI-generated feedback helped individuals learn and improve their skills, it also caused a decrease in intellectual diversity of the population overall. Through a form of algorithmic monoculture, receiving feedback from the same, centralized AI system, individuals tended to specialize in similar ways, thus reducing the intellectual diversity in the population overall. The same study also points to another risk. Although AI systems can enhance collective memory by supporting learning and providing feedback, they may also exacerbate existing inequalities by offering more learning opportunities to those who are already highly skilled and motivated.

These are just two examples of well-intended applications of AI can have unintended consequences that limit the collective memory, attention, and reasoning ability of groups. Well-designed systems that emphasize the humanity of the situation by placing humans in the center while still leveraging AI's capability to deal with large input can enhance collective reasoning by amplifying diverse thinking styles and backgrounds.

To realize the transformative benefits of AI requires thinking beyond the immediate performance and efficiency benefits of automating tasks. Taking a broader view that considers how AI can strengthen collective intelligence by supporting collective memory, attention, and reasoning can open opportunities to unlock the true potential of human-AI collaboration. Applying the lens of collective intelligence can inform both the design and deployment of AI systems in principled ways to increase groups' capacity to solve complex problems and adapt to changing environments.