

People & Organizational Performance Practice

Building expertise in the age of AI: Who trains the next generation?

As AI reshapes entry-level work, organizations should rethink how expertise is developed by integrating knowledge management, role design, learning, and coaching into a single system.

*by Bryan Hancock
with Charlotte Seiler*



For decades, organizations have relied on early-career talent to do the routine, lower-risk work that supports the business and to serve as a training ground for future leaders. Think of Peggy Olson's trajectory on *Mad Men*, from novice assistant to Don Draper's protégé to confident copy chief at an advertising agency—a climb that began because routine work kept her in the room where judgment was practiced, and where her own could be noticed.

In real life, advances in automation and AI are changing the composition of entry-level work itself. Tasks such as research, documentation, data cleanup, basic coding, and preliminary analysis are being streamlined or absorbed into AI systems. These are precisely the activities through which young employees have traditionally built instincts, developed judgment, and earned the right to take on more.

At the same time, fears are growing about the impact of AI on jobs. In the 2025 [Women in the Workplace report](#) by McKinsey and LeanIn.Org, entry-level workers, particularly women, reported feeling the most worried of all age groups about how AI use will affect their jobs. The anxiety is broad based: Among graduating seniors, pessimism about starting a career climbed to 62 percent, from 46 percent, in just two years, and three-quarters of the pessimists pointed to firms hiring fewer entry-level workers as the reason.¹

The outlook for entry-level hiring is still evolving, but early indicators suggest a tightening market. Unemployment among recent US college graduates has increased since 2019, while the number of entry-level roles is declining, particularly in AI-exposed occupations. As of the first quarter of 2026, the unemployment rate for recent college graduates stood at roughly 5.7 percent, and about four in ten were underemployed—working in jobs that do not typically require a degree.² Research from the Stanford Digital Economy Lab finds that workers ages 22 to 25 “in the most AI-exposed occupations have experienced a 16 percent relative decline in employment even after controlling for firm-level shocks.”³

How much of this softening is attributable to AI remains genuinely contested: Federal Reserve Bank of New York economists estimate that the rise of remote work—which makes it harder to train novices at a distance—accounts for much of the increase in young-graduate unemployment, and Yale's Budget Lab finds no clear economy-wide AI fingerprint yet, even as it flags the growing divergence between younger and older graduates as consistent with early-career effects.⁴

Whether the apprenticeship channel is being eroded by algorithms or by distance, the implication for employers is the same: The informal mechanisms that once turned novices into experts can no longer be taken for granted. These shifts point to a narrowing set of traditional entry points just as the nature of early-career work is being redefined.

For organizations, this is a moment of both uncertainty and choice. As AI reshapes how work gets done, the question is no longer how many entry-level roles to hire, but [what those roles are designed to do](#). Companies can use this transition to rethink entry-level work as a foundation for building expertise in an AI-enabled environment—equipping early-career employees to design,

¹ *AI and the workforce ahead: What the class of 2026 tells us about the future of the labor market*, Handshake Network Trends, 2026.

² “The labor market for recent college graduates, Q1 2026 update,” Federal Reserve Bank of New York, May 5, 2026.

³ Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen, “Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence,” Stanford Digital Economy Lab working paper, November 13, 2025.

⁴ Natalia Emanuel, Emma Harrington, and Amanda Pallais, “Remote work leaves younger workers sidelined,” Federal Reserve Bank of New York *Liberty Street Economics*, June 1, 2026; “Tracking the impact of AI on the labor market,” The Budget Lab, updated June 15, 2026.

develop, and steer AI systems, not just operate within them.

This article outlines a four-part approach spanning knowledge management, role design, learning in the flow of work, and managerial coaching. Organizations that move deliberately can accelerate how quickly employees build judgment and take on higher-value problem-solving, strengthening both individual experience and long-term performance.

There will always be a pipeline

Entry-level employees bring fresh perspectives and creative problem-solving to organizations, fueling growth and innovation. They are apprentices to more-seasoned mentors, building relationships that strengthen the culture for everyone. They become middle managers and senior leaders down the road, keeping the pipeline strong into the future.


The question of how to generate the next group of talent is not new. It is, however, being exacerbated as AI use cuts across so many different levels of knowledge work simultaneously. Two senior Microsoft engineering leaders, Mark Russinovich and Scott Hanselman, describe the dynamic bluntly: Agentic coding assistants give senior engineers an “AI boost,” multiplying their throughput, while imposing an “AI drag” on early-career developers who lack the judgment to steer and verify AI output. The resulting incentive—hire seniors, automate juniors—quietly dismantles the bottom of the talent pyramid on which every senior role depends. Their prescription is striking for its candor: Keep hiring early-career employees, accept that they initially reduce capacity, and make their growth an explicit organizational goal.⁵

McKinsey survey data shows that gen AI use is starting to affect the need for entry-level positions at some organizations (Exhibit 1).

Exhibit 1

Gen AI use is starting to affect entry-level roles at many organizations.

Gen AI’s impact on need for entry-level roles,¹ % of respondents



Note: Figures do not sum to 100%, because of rounding.
¹Question: Is gen AI reducing the need for entry-level roles in your organization?
Source: McKinsey’s New Era of Work Survey, Sept 2025 (n = 28,000)

McKinsey & Company

⁵Mark Russinovich and Scott Hanselman, “Redefining the software engineering profession for AI,” *Communications of the ACM* April 2026, Volume 69, Number 4.

While that is a trend worth monitoring, leaders we speak with are committed to hiring and apprenticing the next generation of experts, even if their numbers are lower than before. And some organizations are holding their numbers steady. Bank of America, for instance, is bringing in nearly 4,000 summer interns and full-time campus recruits in 2026—matching last year's hiring from a talent pool spanning more than 500 schools—while explicitly redesigning those roles around AI from day one. Its chief people officer calls the approach “intentional and long term.”⁶

In a March 2026 survey of roughly 1,500 executives and senior talent leaders, those expecting AI to increase entry-level hiring in 2026 outnumbered those expecting it to decrease hiring by nearly three to one—though the same survey carries a warning, with a third of employers reporting that AI has already reduced the foundational, skill-building tasks juniors learn from.⁷

Organizations want to hire people who have the nascent skills to interpret the in-depth data generated by AI and who can work alongside agents. To create a strategy that guides these employees, leaders can focus on four areas.

The foundation: Knowledge management

Organizations must leverage their best knowledge to develop AI workflows. If they don't know how to marshal their current expertise to create high-quality data, it becomes difficult to foster the next group of experts. The goal is to make expertise usable, reliable, and embedded in day-to-day workflows.

Codified expertise captures how top performers think—their frameworks, decision rules, assumptions, and past judgments—and structures that knowledge so large language models (LLMs) can access it. This includes clean data, curated case materials, and standardized taxonomies.

For example, consider a junior associate at a law firm working on an M&A contract. Rather than simply searching across thousands of past deals, an AI-enabled system can surface the most relevant precedents based on deal characteristics and highlight the clauses, trade-offs, and negotiation patterns used by top partners in similar situations. In doing so, it not only accelerates drafting but also exposes the reasoning behind key decisions—helping the associate build judgment while doing the work.

A similar challenge arises in determining which knowledge should be elevated. A third-year employee at an engineering firm may document a single client engagement clearly and accurately. But a senior leader, drawing on decades of experience across many projects, may produce analyses that reflect broader patterns, trade-offs, and long-term implications. Both inputs are valuable, but they are not equivalent. For knowledge management to be effective in an AI-enabled environment, organizations need ways to distinguish between isolated experience and accumulated judgment.

That distinction should be explicit in how knowledge is curated and surfaced. Systems can

⁶“BoFA to welcome nearly 4,000 summer interns and campus recruits,” Bank of America press release, June 3, 2026; Katherine Doherty, “Bank of America bucks AI job fears with 2,000 summer interns,” Bloomberg, June 3, 2026.

⁷Andrew R. Hanson and Molly Cook Escobar, *Entry-level hiring in the AI era: What employers are thinking (and doing)*, Strada Education Foundation, May 19, 2026.

weight inputs based on factors such as depth of experience, repeatability of outcomes, and relevance across contexts. This ensures that AI tools and junior employees are guided not just by what is documented, but by what reflects the strongest underlying expertise.

At its core, modern knowledge management indexes and prioritizes information in ways that reflect how expertise is built, including expert validation, feedback loops, and clear ownership. The system improves as it is used, capturing new insights and refining what “good” looks like. This enables less experienced employees to build judgment faster and take on more complex problem-solving earlier in their careers.

Entry-level role design: The ‘answer key’ model

As AI and agentic workflows take on more execution of day-to-day business, entry-level roles can be redesigned around working with, supervising, and improving AI-driven outputs. This creates an opportunity to augment near-term productivity while preserving the long-term pipeline.

In an AI-enabled environment, the goal is no longer just task proficiency, but judgment: knowing when to trust, question, or override machine-generated outputs. When knowledge management and role design are strong, early-career employees can access institutional expertise far earlier. The limiting factor becomes their ability to apply that knowledge in context.

As Matt Beane argues in *The Skill Code*,⁸ early workers’ skills develop through a combination of challenge, complexity, and connection to expert thinking through apprenticeship—conditions that are often reduced when technology absorbs routine work.

To avoid this disconnect, entry-level roles can include a more intentional form of apprenticeship, with AI in the loop. Teams can create roles for lower-risk contributions while building in frequent coaching on how to interpret and refine AI outputs. Some organizations are introducing curriculum-based roles, where employees progress through defined assignments that pair human work with AI assistance. Others are creating parallel workflows, where entry-level employees complete work independently and then compare it against AI-generated results, using the gap as a learning tool.

Call this the answer-key model: The employee attempts first, the AI grades the attempt, and the employee sits down with a manager, who discusses the differences. In this model, AI does not replace apprenticeship but reshapes it, accelerating feedback while preserving the conditions required to develop expertise.

At one real estate firm, an AI agent generated highly detailed market assessments that were more comprehensive than what an entry-level employee could produce independently. Rather than relying on the tool alone, the company asked junior employees to build their own assessments “by hand,” walking neighborhoods, studying geography and traffic patterns, and forming an independent view. They then compared their analysis with the AI-generated output. The contrast served as a structured learning tool by highlighting gaps, surfacing missed factors,

⁸ Matt Beane, *The Skill Code: How to Save Human Ability in an Age of Intelligent Machines*, Harper Business, 2024.

and reinforcing strong judgment. In this model, the agent accelerates feedback and broadens perspective while the employee continues to develop the contextual understanding and instincts that underpin expertise. Used this way, AI sharpens entry-level learning.

Evidence from other fields suggests the comparison step is what does the teaching. In clinical studies, simply giving physicians an LLM barely improved their long-term diagnostic performance; a workflow that required them to compare and reconcile their own reasoning with that of the AI model lifted performance in future situations to the level of the AI model alone.⁹

The inverse result is just as instructive. When workers used gen AI to perform technical tasks they could not do themselves, the capability vanished the moment AI access was removed. No durable skill had formed. Passive reliance builds output; structured comparison builds experts. One caution: Workflows that demand this kind of cognitive engagement are harder to use and can overwhelm beginners, so they must be deliberately designed and dosed, not bolted on.¹⁰

To manage risk, many organizations are expanding sandboxed environments for early-career talent. Bank of America gives interns AI training specific to their line of business from day one and is using simulation to compress the judgment building that junior bankers once accumulated through routine work.

“We’ve got to give people the experiences in a simulated way quickly,” says Josh Bronstein, the bank’s head of global talent, “so that they have the judgment that they otherwise would have gained by doing some of those tasks in their first year or two, that now, or soon, AI or other technology can do.”¹¹ These employees can practice in internal simulations or in external contexts, such as working with not-for-profit organizations, before applying these skills in higher-impact settings. This is increasingly important as junior employees begin to configure or direct AI agents that can scale decisions quickly.

The implication is that entry-level roles are becoming more about learning how work gets done in an AI-augmented system—developing the judgment, oversight, and adaptability required to operate effectively as both a contributor and supervisor of digital labor.

At the same time, the structure of entry-level roles themselves is shifting. Rather than confining new hires to narrow, task-based specialties, some organizations are redesigning roles around end-to-end system thinking and broader responsibilities.¹² Instead of hiring a sales analyst, for instance, a company may hire for a broader role that encompasses sales, marketing, and commercial expertise. Since generalists have the tools to customize for a specific problem or context, work can be designed more for them.

Don’t wait: Learning in the flow of work

Organizations can create agents and learning systems that extend their deep knowledge and track skill development in real time. The answer-key model even supplies the metric: The gap between an employee’s independent attempt and the AI model’s output is observable, and a gap that narrows over time is direct evidence that judgment is forming—something apprenticeship,

⁹ Goh et al., “Large language model influence on diagnostic reasoning,” *JAMA Network Open*, October 2024, Volume 7, Number 10; Goh et al., “GPT-4 assistance for improvement of physician performance on patient care tasks: A randomized controlled trial,” *Nature Medicine*, February 2025, Volume 31.

¹⁰ Jenna Butler et al., *Microsoft new future of work report 2025*, Microsoft, December 2025.

¹¹ Alice Tecotzky, “Bank of America accepted less than 1% of internship applicants this year, but AI isn’t shrinking the class,” *Business Insider*, June 3, 2026.

¹² Aili McConnon, “The bottom rung returns as AI reshapes entry-level jobs,” *IBM*, March 2, 2026.

for all its history, has never been able to measure. The idea of “learning in the flow of work” becomes more concrete in this context because entry-level employees learn while producing real outputs—with AI systems acting as embedded guides and reviewers.

As companies redesign workflows to include AI agents, much of the routine execution is handled by the system. What remains for humans, including junior employees, is higher value: interpreting outputs, making trade-offs, and connecting insights across contexts. The nature of productivity changes as entry-level workers contribute at a higher level earlier in their careers. For this to happen, continuous support is built into the work itself. AI tools can suggest next steps, flag risks, and provide instant feedback, while managers focus their coaching on judgment and decision-making rather than task mechanics.

In practice, some organizations front-load learning through short simulations or structured onboarding. Others move quickly into real work, trusting that AI-enabled feedback loops and coaching will accelerate development in real time. In both cases, the model is the same: Employees are not waiting to become productive; they are learning by doing, with AI compressing the time it takes to build experience and earn trust.

The learning science behind the answer-key model is beginning to firm up. In one experiment, first-year medical students who answered AI-generated case questions and received automated, personalized feedback over five days outperformed second-year students—a full year ahead in training—on the targeted diagnoses in a video-based exam, including at a two-week follow-up. The learning persisted, providing evidence that the attempt-then-check loop can produce durable skill rather than borrowed competence.¹³ A Stanford pilot pushes the design further: A chatbot plays the patient, and the student interviews the patient, commits to a diagnosis, and defends it, and only then does the system critique the reasoning.¹⁴ The sequencing is the key point: The student goes first, and AI grades the attempt.

Hire for judgment, coach for context

Organizations are thinking differently about who they should hire. Instead of building large cohorts of junior specialists, they are placing greater emphasis on adaptability over narrow expertise. The question is shifting from “What did you study?” to “How do you think?”

In addition to digital literacy, employers are prioritizing attributes such as creativity, problem-solving, resilience, and ability to reason—qualities associated with high general cognitive ability.¹⁵ This pattern shows up in hiring data: When employers rate the skills that matter most for entry-level hires, these skills are at the top of the list (Exhibit 2).

The durable bet, in other words, is judgment, not tool fluency. Roles that demand AI skills are also nearly twice as likely to demand analytical thinking, resilience, or agility alongside them.¹⁶ One head of HR at a Fortune 100 financial firm described hiring not for a specific degree, but for “general athletes” with strong learning instincts, relational skills, and baseline familiarity with AI.

This shift also opens the door to a broader set of workers, including what the not-for-profit

¹³Yavuz Selim Kiyak et al., “AI teaches surgical diagnostic reasoning to medical students: Evidence from an experiment using a fully automated, low-cost feedback system,” *Journal of Surgical Education*, October 2025, Volume 82, Number 10.

¹⁴Carrie Spector, “Stanford education scholar uses AI to help medical students hone diagnostic skills,” Stanford Graduate School of Education, September 30, 2024.


¹⁵Andrew R. Hanson and Molly Cook Escobar, *Entry-level hiring in the AI era: What employers are thinking (and doing)*, Strada Education Foundation, May 19, 2026.

¹⁶Elina Mäkelä and Fabian Stephany, “Complement or substitute? How AI increases the demand for human skills,” arXiv, revised June 8, 2026, arXiv:2412.19754.

Exhibit 2

Organizations are prioritizing skills associated with high general cognitive ability over task-based capabilities.

Top future skills according to HR professionals, 2025–26, % of respondents (who mentioned the skill among their top 5)



Note: Year-over-year differences should be interpreted as directional only, as country coverage changed slightly compared with 2025, with the addition of Belgium, China, and the Netherlands.
Source: McKinsey HR Monitor Survey, Jan 2026, n = 1,303 HR professionals in Belgium, China, France, Germany, Italy, Netherlands, Poland, Spain, UK, and US; McKinsey HR Monitor Survey, Dec 2024, n = 1,925 HR professionals in France, Germany, Italy, Poland, Spain, UK, and US

McKinsey & Company

organization Opportunity@Work calls “STARs” (Skilled Through Alternative Routes)—individuals without four-year degrees who have built capabilities through experience. Because knowledge is more accessible, these workers can contribute meaningfully if given the right support and development opportunities.

This marks a departure from the traditional hiring model that rewarded deep specialization early on. As AI and knowledge systems increasingly provide on-demand expertise, the value of employees lies in how effectively they apply, connect, and extend that knowledge. In practice, companies are selecting more well-rounded, high-potential individuals—and relying on apprenticeship and real-world experience to build domain expertise over time.

Because junior employees are being pushed into higher-value activities much earlier, coaching

for human skills is even more important. A new hire may be in the position of engaging with senior stakeholders sooner than they would have in the past. They may have to deliver difficult messages, such as recommending a significant budget reduction, and do so with judgment, tact, and confidence. Managers can teach a new analyst how to navigate a sensitive conversation, build credibility, or influence an experienced operator.

Without guidance, early-tenure employees may default to presenting “what the data says” in a way that feels blunt or disconnected from business reality. Managers have always coached younger employees to ask thoughtful questions and incorporate other perspectives. But these traits are even more important today, since these employees can interpret more potent and comprehensive data than in the past. Few senior executives have ever wanted to be told by a newly minted MBA that their business plan is unrealistic. Junior employees may be backed by AI, but they still need to be coached to read the room.

At the same time, coaching can focus on helping new hires build contextual understanding. As analysis becomes increasingly automated, the differentiator is not producing the number but explaining it. New hires need to learn how to interpret patterns, recognize industry dynamics, and articulate the “why” behind the data. This kind of judgment typically comes from experience, but in this new environment, coaches can accelerate that learning through real-time feedback, storytelling, and embedded “micro lessons” that explain how the business actually works. Helping an analyst understand, for example, that a spike in sales may be driven by promotional cycles rather than true demand is just as important as teaching them how to present the numbers.

Some suggest taking inspiration from the medical profession on formalizing coaching in the flow of work, giving it the institutional weight that informal mentorship has lost. Molly Kinder, an expert on AI’s impact on work, suggests that employers borrow from the medical residency model. If AI absorbs the routine tasks through which juniors once learned incidentally, companies should redesign entry-level roles as protected, structured periods of deliberate skill building, with progression toward independence as an explicit commitment.¹⁷

Russinovich and Hanselman take that logic one level deeper into the daily workflow, proposing a “preceptor” model—a clinical arrangement in which a new nurse or physician practices under a designated senior before earning the right to work independently.¹⁸ In their version, a senior engineer formally mentors a small cohort of juniors, working with AI tools together so the senior can observe what the junior accepts, rejects, and misjudges—shifting the mentor’s job from answering questions to teaching judgment.

Ultimately, organizations should treat coaching as a core capability in developing early-career talent, not as an afterthought. Managers are no longer simply supervising work; they are responsible for shaping how employees think, communicate, and engage with the business. This may be reinforced through more structured coaching, better knowledge-sharing systems, and redesigned entry-level programs that provide broader exposure to different functions. As entry-level roles evolve from executing analysis to advising the business, success will depend on how

¹⁷ Molly Kinder, “To save entry-level jobs from AI, look to the medical residency model,” Brookings Institution, January 23, 2026.

¹⁸ Mark Russinovich and Scott Hanselman, “Redefining the software engineering profession for AI,” *Communications of the ACM* April 2026, Volume 69, Number 4.

effectively coaches can equip new hires with the human skills of judgment, communication, and influence that are required to operate in this new reality.

What early-career employees want from work hasn't changed

Early-career workers are facing fewer opportunities and slower mobility, yet their expectations for growth have not shifted—creating a gap that organizations must actively manage rather than assuming it will resolve itself. What employees want from work has remained largely unchanged: meaningful roles, development, flexibility, and strong leadership.

In this environment, retraining and retention are tightly linked. Retention requires delivering on the fundamentals, including clear career paths, meaningful work, supportive managers, and access to in-demand AI skills. This is especially critical because the employees most engaged with AI are also the most likely to leave, given their market value.

The implication is straightforward. Companies cannot treat early-career talent as a transient workforce. With mobility down, there is a window to invest more deeply, but also a risk of disengagement if expectations are not met. The organizations that succeed will be those that pair intentional skill building with a compelling employee experience, ensuring entry-level workers not only adapt to AI-driven change but also choose to stay and grow within it.

The investment case answers the obvious objection. If junior employees initially reduce capacity, why should any one firm bear the cost of building judgment that a competitor might poach—particularly when AI-engaged employees are the most mobile? While firms have always faced the risk of [training talent](#) for competitors, the difference now is that AI is leading some to question the value of maintaining entry-level pathways.

Part of the answer is that the training described in this article is not generic. Judgment built based on a firm's own codified expertise, cases, and taxonomies is partly firm specific; it travels less easily than a certification. There is also the cost of abstaining: weakened succession, stalled knowledge transfer, and slower AI adoption itself. These represent the longer-term risks that leaders at firms investing through this transition cite as the reason a juniors-versus-efficiency trade-off is a false binary.

Although uncertainty abounds about how much AI will reshape the workforce, it's worth remembering that the pipeline is not disappearing. It is being rebuilt. Organizations that redesign roles to steer AI toward productive outcomes will be better positioned to develop the skilled workforce they need.

The need for deliberate design

Realizing the potential of early-career employees is not automatic. It requires intentional design in four areas:

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



- *Knowledge management.* Codify how your best performers think—their frameworks, decision rules, and past judgments—and structure knowledge so AI systems and junior employees can draw on it directly. Build in expert validation, feedback loops, and clear ownership, and weight inputs so accumulated judgment outranks isolated experience. The system should improve as it is used and continually refining what good looks like.
- *Workflow and role redesign.* As you redesign workflows around AI agents, treat entry-level roles as a deliberate part of that redesign rather than a residual. Rebuild those roles around working with, supervising, and improving AI outputs. Hire for judgment over tool fluency—curiosity, adaptability, problem-solving, and emotional intelligence—and widen the aperture to include strong candidates without formal training in the given field. The goal is to hire a “general athlete” who can apply, connect, and extend knowledge across a workflow, not a narrow specialist.
- *AI-based learning design.* Embed learning into real work so employees build judgment while producing genuine outputs, with AI acting as an embedded guide and reviewer. Anchor the design in the attempt-then-check loop, and treat the narrowing gap between an employee’s independent work and the AI model’s output as direct evidence that judgment is forming. Design and dose these workflows deliberately rather than bolting them on.
- *Manager upskilling.* Treat coaching as a core capability, not an afterthought, since junior employees now engage senior stakeholders and interpret potent data far earlier than before. Equip managers to coach judgment, context, and influence rather than supervising task mechanics. Consider formalizing the role through a preceptor model, giving structured mentorship the institutional weight that informal apprenticeship has lost.

Organizations that continue investing in early-career talent while redesigning roles to work alongside AI are more likely to develop the skilled workforce they’ll need in the future. Indeed, one of the most notable findings of the early AI era may be how little the hierarchy of essential skills has changed, even as the bar for entry-level work rises.

Bryan Hancock is a partner in McKinsey’s Washington, DC, office, and **Charlotte Seiler** is an associate partner in the Bay Area office.

This article was edited by Barbara Tierney, a senior editor in the New York office.

Copyright © 2026 McKinsey & Company. All rights reserved.