

ARTICLE

Does AI Require Rethinking Education?

A perspective on what AI means to universities and employers' on how to approach the next generation

My previous article was about what AI is doing to the consulting industry, which traditionally skews toward seniority (and probably will continue to). But there's another side to the AI disruption that deserves just as much attention: what happens to the next generation of graduates when AI takes over precisely the tasks juniors have traditionally learned the trade on? That question is also being addressed today by [the director of Denmark's central bank, Nationalbanken](#).

Senior consultants have an edge: they've built up experience with the organizational decision-making processes and the inertia that often gets in the way of value creation. Seniors typically have the experience it takes to correct AI and develop genuinely good results with it.

New graduates and junior consultants appear to be finding it harder to deliver value that matches the hourly rate their work is sold at – so how should education adapt to this new reality?

I set out to look at what various thought leaders have to say about the challenges juniors face, and condensed the answers into this piece.

Forward-looking trends

Entry-level work isn't disappearing – it's being redefined

The number of entry-level job postings in the US has fallen by 35 percent over the past 18 months, largely because of AI, according to figures cited by [World Economic Forum](#). But that doesn't mean the junior role is disappearing – it means it needs a clearer answer to which tasks it should be solving.

Organizational units will still be built around a staff of a few seniors and more juniors – not because juniors are cheaper to put on routine tasks, but because a large volume of output is generated that seniors have to act on. AI doesn't reduce the total amount of work that requires human judgment and accountability – it shifts it. An [analysis of AI's impact on](#)

[finance functions](#) describes exactly this mechanism: the more AI produces, the more still needs to be validated and approved, and that capacity doesn't simply grow because AI does.

It's now worth focusing junior development on the ability to work independently within a direction set by seniors. No leadership will let a new graduate sign off on AI-driven decisions across the organization. But they can prepare the decision brief: analyze, assess AI's output, and come up with a recommendation that a senior then decides on. How large a scope they're allowed to prepare on their own depends on how mature they are – precisely the model a [Forbes analysis on redesigning entry-level tech roles](#) describes as a “reversibility budget”: bounded, reversible authority that grows in step with demonstrated judgment.

[Harvard Business Review](#) describes the same logic from another angle: instead of giving juniors responsibility on the big, risky engagements, more industries are giving them real, but bounded, autonomy on smaller, standardized tasks early on.

That points straight ahead to another question: how much of that independence and judgment can education build in before someone even gets their first job? I'll come back to that.

Traditional competencies don't predict who succeeds with AI

A study by KPMG and the University of Texas involving 523 early-career professionals, covered in [Harvard Business Review](#), sorted juniors into three groups based on how they work with AI:

AI apprentices

AI delegators

AI amplifiers

The surprising point is that classic indicators of strong theoretical knowledge couldn't predict who ended up in which group. What actually set them apart was something else: how they use their acquired skills together with AI and get AI to amplify them.

‘Apprentices’ struggle to bring their own strong skills into play and often send AI in the wrong direction.

‘Delegators’ are productive and appear good at leveraging AI, but also trust its output without critical reflection.

‘Amplifiers’ are strong at framing problems in ways that guide AI in the right direction, while iteratively refining the result.

How a new approach can take shape

The optimistic message is that the *apprenticeship model* can be an inspiration for a new approach. That requires reshaping it – from observation, where the trainee becomes able to repeat what the mentor does, into a more equal partnership between trainee and mentor.

Make pairing with an experienced senior a formal structure – not a coincidence.

Students and new graduates should be systematically paired with one or more experienced seniors, specifically to accelerate their ability to exercise “AI discernment” – that is, recognizing when an AI output is usable, and when it should be challenged.

An analysis in [Forbes](#) on what the author calls “intelligent talent density” points to the same model: juniors contribute AI fluency and a fresh perspective, while seniors validate output against real-world constraints – in practice, “reverse mentoring” upward combined with joint validation downward. It's an apprenticeship built around assessing AI work together, instead of observing and copying a senior.

Teach juniors to close the gap on seniors' head start. Give juniors a basic understanding of the incentive structures behind a client request – financial gain, political motivation, organizational dynamics – and how these translate into what the client is actually asking for. That will make juniors far better able to put AI to use. As a complement to seniors, juniors actually have something of an advantage in not being weighed down by too many preconceptions. That pushes juniors toward fresh thinking and modernization, just as juniors can find it easier to adopt new technologies.

AI is, in itself, an enormous knowledge base. Juniors should actively use an alternative AI as a second opinion, to challenge and test what the first AI came up with, instead of simply accepting the first answer. The practice is already widespread in AI-heavy organizations. [McKinsey](#) describes “agent-as-judge” as a standard part of evaluating agentic AI. But it isn't a risk-free shortcut: [MIT Sloan](#) warns that using AI to check AI can become a “tempting shortcut” that reinforces shared blind spots rather than exposing them – which is exactly why juniors need to be trained to judge for themselves, not just outsource the judgment to another AI.

AI discernment will become the decisive factor in whether tomorrow's employee succeeds or fails. The decisive competency isn't knowing the most – it's being able to identify when AI's output is usable, and when it should be challenged. Theoretical knowledge isn't enough – it has to be trained.

Make “AI discernment” a discipline in its own right – not a side effect of enough years of experience. Today, that ability arises as a byproduct of many years of practice. If the KPMG/UT study is right that classic experience isn't the decisive factor, it should be possible to teach it more directly: concrete exercises in stress-testing AI output, learning typical failure patterns, and practicing recognizing “this looks right, but it isn't” – rather than waiting for it to emerge on its own after ten years in the industry.

What should education and employers actually do?

The KPMG/UT study doesn't just offer a diagnosis – it also gives a concrete answer for what actually works. It's worth translating into two separate areas of responsibility.

Education should teach how the work is actually done – not just the theory behind it.

If strong theoretical knowledge alone doesn't make the difference, as the study shows, teaching should be built around simulated client assignments. Students need to practice the craft itself: steering AI workflows, challenging output, and iteratively refining the result, the way the study's ‘amplifiers’ do it in practice.

Employers should make judgment visible and train the path to it – not just the final answer. The study's authors point to two concrete moves themselves: explicitly document why an AI output was accepted, changed, or rejected, so the reasoning becomes visible and trainable – instead of staying a silent instinct. And assess employees on how they interact with AI along the way – whether they question, challenge, and justify – not just on the finished product, which AI can already deliver on its own.

This isn't just theory. [KPMG](#) has built its own internal AI onboarding program around exactly this: personalized learning paths, scenario-based exercises that mimic real client work, and an internal network of “AI champions.” That, by all appearances, is what separates the organizations that manage to turn AI into real value from those that don't.

Theoretical Knowledge Isn't Automatically Good Business

It isn't only consulting firms and companies that need to reinvent themselves. There's already evidence that education can be redesigned to make students genuinely job-ready – and that most institutions haven't done it yet.

Aalborg University has run its teaching on a [Problem Based Learning model](#) since 1974, in which students “learn best by actively applying research-based theories and methods to real problems.” Instead of first receiving theory and then applying it, students are required to “identify, formulate, and address a current, relevant problem in its context.”

The employer side shows how wide the gap still is almost everywhere else. An [employer survey from Poets&Quants](#) finds that 77 percent of employers expect new hires to have experience with AI tools, while 58 percent believe universities aren't doing enough to build those skills – and only 3 percent of alumni took a university course in AI at all; seven out of ten instead taught themselves by experimenting on their own.

An article from [Ohio State University](#) points to what the answer should be: universities need to move from being academic institutions to actually functioning as career-preparation pipelines – structurally embedded internships for all students, AI competency built into every field on par with a core subject like mathematics rather than treated as an elective. But perhaps the most important point is that universities' career centers need to be elevated from peripheral student services to a core academic function.

Theoretical knowledge isn't automatically good business. The decisive competency isn't knowing something – it's being able to turn what you know into value in a specific business context. You can't read your way there; it has to be practiced.

Aalborg University shows that it can be done at scale, and that it can sustain an entire university for half a century. The question is no longer whether the model works – it's why it's still the exception rather than the norm.

Closing considerations

If AI's arrival removes the tasks juniors have traditionally learned the trade on, without replacing them with something new, the industry loses the talent it will depend on ten years from now. But the solution, in my view, isn't to slow down AI or defend the status quo – it's to make the apprenticeship model something we actively design, instead of something that simply emerged as a byproduct of how work used to be organized.

The common thread, fundamentally, is that education and industry need to move closer together. Vocationally oriented programs – the ones that, like Aalborg's PBL model, are already built around real problems and close contact with employers – are markedly better equipped to prepare juniors for an AI-heavy working life than the classic faculties, where change has traditionally moved more slowly.

And then a tongue-in-cheek disclaimer, aimed at this very piece: as career author Lindsey Pollak put it in [an article from Ohio State University](#), “we're terrible at predicting which majors or skills will matter next” – an entire generation was steered toward software

development, only to watch many entry-level programming jobs disappear just as they graduated. So take these suggestions too, on what juniors should be able to do, with a grain of salt.

Sources

1. [Research: Why Some Junior Employees Work Well with AI \(and Others Don't\)](#) — Harvard Business Review, July 2026
2. [How AI is changing the nature of entry-level work](#) — World Economic Forum, March 2026
3. [Intelligent Talent Density: Synthesizing Junior Agility And Senior Context](#) — Forbes, September 2026
4. [Redesigning Entry-Level Tech Roles To Build Better Judgment](#) — Forbes, September 2026
5. [How AI Is Upending How Consulting Firms Hire Talent](#) — Harvard Business Review, October 2025
6. [AI Reduced The Work, But Human Oversight Became The Bottleneck](#) — Airwallex
7. [Employer Survey: Universities Aren't Producing Enough AI-Ready Graduates](#) — Poets&Quants, 2026
8. [What Is College for in the Age of AI?](#) — Ohio State University, Office of Academic Affairs, January 2026
9. [Problem Based Learning \(PBL\) at Aalborg University](#) — Aalborg University
10. [Evaluations for the Agentic World](#) — McKinsey/QuantumBlack, January 2026
11. [Seeing Real Value from AI Depends on Being Able to Verify Its Outputs](#) — MIT Sloan, June 2026
12. [AI Is Changing the Labor Market – Two Groups Are Particularly Vulnerable, Says Nationalbanken](#) — TV 2, September 2026