

The federal government's new AI strategy is, at its core, a bet on human capital. Train a million Canadians. Equip every post-secondary student with AI tools. Build a workforce that can compete. It's an ambitious vision — and one that universities are already scrambling to meet with workshops, seminars, and hastily assembled ethics guidelines.

But here's the uncomfortable truth: literacy isn't enough anymore.

When we talk about AI literacy, we mean something fairly modest — the ability to understand what AI is, engage with it critically, and use it responsibly. These are important foundations. But the landscape has shifted beneath us faster than most institutions anticipated. The question is no longer whether students can *use* AI. It's whether they can *think with it, build with it, and govern it*.

That requires something closer to fluency.

Think of the distinction this way: a literate person can read a newspaper. A fluent person can write one, fact-check it, argue with its editors, and launch a competing publication. The gap between those two capacities is enormous — and it's precisely the gap that most AI education programs are currently failing to bridge.

The rise of agentic AI makes this gap urgent. Today's most powerful AI systems don't just answer questions; they plan, execute multi-step tasks, use external tools, and operate as nodes within larger automated workflows. Google's Co-Scientist, for instance, uses a multi-agent architecture to generate and stress-test scientific hypotheses autonomously. Research tools like Elicit and Undermind are surfacing cross-disciplinary connections that no single human researcher could reliably find. These aren't novelties — they're the infrastructure of tomorrow's knowledge work.

A student who understands AI conceptually but can't engage dynamically with these systems

— integrating them into research, evaluating their outputs critically, designing workflows around them — will be at a profound disadvantage. Only 14 percent of current graduates report strong proficiency in applying AI tools to professional work, according to a recent survey by Pearson and Amazon Web Services spanning over 2,700 respondents across six countries. Meanwhile, more than half of employers say they struggle to find AI-ready candidates.

Universities have an opportunity — and, arguably, an obligation — to close that gap. But doing so requires rethinking what AI education actually is.

True AI fluency education should go beyond tool familiarity. It means cultivating the judgment to know when to trust an AI output and when to interrogate it. It means understanding data provenance, algorithmic bias, and the ethics of automated decision-making at a level of depth that goes well beyond a one-hour workshop. It means training students not just to consume AI systems but to contribute to their design, evaluation, and governance.

This has implications that extend well beyond workforce preparation — though that matters too, with Canadian job listings requiring AI skills climbing steadily through 2025. Democratic institutions, scientific integrity, public health systems, environmental policy: all of these depend on people who can hold AI accountable, not just people who can prompt it.

Several universities in North America have adopted AI literacy frameworks — from Stanford's model to EDUCAUSE's higher education guidelines — but most of these frameworks stop short of the fluency threshold. They offer taxonomies of competencies without grappling seriously with the subject-specific, domain-embedded nature of real AI integration. A nursing student, a climate scientist, and a policy analyst all need AI fluency, but they need it to look quite different.

