

Use it or Lose it - AI and the Atrophy of Skill

If we automate away the work that builds expertise, where does the next generation of experts come from?

- [Use it or Lose it - AI and the Atrophy of Skill](#)

Use it or Lose it - AI and the Atrophy of Skill

Forward

Years ago I managed a small team of talented developers where writing SQL was a large part of the job. The team's SQL experience varied, some were early in building it and some were fairly advanced, and part of my role was reviewing queries they were wrestling with. Sometimes a developer would bring me a query they used AI to generate. I could usually tell. The logic worked, mostly, but there was something about the structure that was subtly *off*. By this time I had spent several years in roles as a DBA and front-end developer, and written thousands of queries, so whatever was off tended to stand out to me.

The part that stayed with me is *why* I could see it. It certainly was not because I was a better developer, the team was very talented, quite frankly better than I am or likely ever will be. It was because I had years of reps they hadn't had yet, and those reps are exactly what teaches you to notice when something is subtly wrong. The tool let them produce a query without doing the reps, which is the gap this essay is about. The difference in experience that lets you evaluate what a tool produces, and leaning on the tool before you have built up that skillset.

Supplement not substitute

This is not an anti-AI piece. I use AI frequently, and I think it is one of the most genuinely useful tools to arrive in my working life. The argument here is narrower, and more important than "AI good" or "AI bad."

AI should supplement a skill, not be a substitute for building it. In the moment, using it looks the same either way. Over time, the two diverge greatly. The danger is not the tool, it is being reliant on the tool before you have the competence it stands in for. The loud voices on both sides miss this distinction. The boosters say AI makes everyone better, the doomers say it makes everyone worse. Both are wrong in the same way, because the truth is, there is a tradeoff. What you gain in speed you can lose in skill.

Skills atrophy and the pipeline

Any skill, unused, decays. We all already understand this, and see it when it is a language we have not spoken in a decade or an instrument we stopped playing. It is equally true of the skills we are now handing to AI.

My own ability as a programmer has measurably declined since I started using AI to help write code. Not code I couldn't write, the simple programs I have written are not difficult, but code I used to be able to write fluently I now find myself reaching for help with out of habit. It is so much easier, and frankly more satisfying, to implement a new feature set and ship it to production in a single afternoon than to spend several days typing it all out by hand. I make that trade with my eyes open, which is why I am not writing this from any position of superiority. I live inside the problem I am describing.

The individual decline is the small version. The large version is institutional, and worries me more. If AI replaces the entry-level work, the queries a junior writes to learn, the code a beginner grinds through to understand how to develop software, then we shrink the population that ever builds the foundational skill in the first place. The experts of tomorrow come from the juniors of today. There is no senior practitioner who was not first a bad one, then a mediocre one, then a competent one; each stage built on the previous. Remove the bottom rungs of that ladder and you do not just lose the juniors. You lose the seniors they would have become, because the pipeline that produces expertise runs through the same work we are most eager to automate away. And what that work builds, more than the ability to do the task, is to judge it.

You cannot catch what you cannot evaluate

When my car breaks down, I take it to a mechanic. That means I know *how to get a car fixed*. It does not mean I know *how to fix a car*. This is a perfectly fine arrangement up until the moment I am stranded somewhere with no mechanic or until a dishonest one tells me I need a repair and I have no way to know whether they are right. Outsourcing a skill is fine when you will never need it yourself and never need to judge whether it was done well. The trouble is that we are outsourcing skills we very much do need, and losing in the process, the ability to tell whether the job was done right.

AI hallucinates: it produces confident, plausible, wrong answers. It carries the biases of its training. These are the common criticisms, and they are true, but they miss a more important point.

The problem is not only that AI is sometimes wrong. It's that a person who never built the underlying skill cannot tell when it is wrong. A developer who can only "vibe code," without a mental model of "good," might ship flawed code if nothing in their experience flags the flaws. This error-checking happens in the practitioner's own head, and that faculty is built the slow way; by doing the work until you know in your bones what "good" looks like.

This is why the atrophy matters so much more than it first appears. The skill you let erode is not just the skill to *do* the work. It is the skill to *evaluate* the work, which is the exact faculty that makes AI safe to use. Lose it, and you are trusting the output of a system you have no way to validate.

Homogenization, and the median of what

There is a further cost, and it took me a while to phrase it precisely. Early on, company leadership told us that using AI would make our work better, and that it will get better the more we use it. My question at the time was better *how*? What data was it trained on, and who decided it was better than what we, a community of Identity and Access Management practitioners, could produce ourselves? How does it produce something better than what it was trained on, if we do not supply it with better? And if it learns from the corrections we feed back to it, whose corrections does it trust? What stops us from teaching it our mistakes? I was, I later learned, fumbling toward a phenomenon research has since named "Model Collapse." Models trained on their own output, or on a flood of mediocre input, degrade rather than improve.

But what defines "mediocre?" I've heard it said that AI is above median human capability. My question is median of what population? Median of all people in a particular skill? It clears that easily. Median of entry-level practitioners in a skill? It probably clears that too. Median of a genuine expert? Not a chance, outside narrow and specialized models. But when a whole field leans on the same tool and calls the result "good enough," the effect is not that everyone becomes average. It is that everyone converges toward the same competent-but-unexceptional middle. What erodes is the work that never came from the median of anything: the outliers, the idiosyncratic, the strange yet excellent. Exceptional work has never been the median of what came before.

The calculator and the crutch

A scientist reaches for a calculator without a second thought, and no one thinks less of them for it. The reason is simple. They already know how to do the math. They could work it out by hand, they know what a right answer looks like, and they would catch the calculator if it ever lied to them. The tool makes them faster. It does not make them dependent, because the competence came first and the tool sits on top of it. Take the calculator away and the scientist is slower, not helpless.

That is the distinction. A calculator is a tool to the person who has the skill and a crutch to the person who does not. The identical calculation, made by someone who could not have understood it well enough to judge it, is a crutch. The difference between the two is invisible in the moment and decisive over a career.

So I am not going to tell you to stop using AI. I use it, and I will keep taking the days it gives me back in place of one afternoon. But I am going to keep being a little uneasy about the skill I can feel it costing me, because that skill is what keeps AI a calculator and not a crutch. Use it to go faster at what you already know how to do. Be very careful about using it in place of learning to do that thing yourself. And guard, above everything else, the one skill this whole essay has been circling: the judgment to know when the confident machine is confidently wrong.