



MASTER THE CRAFT

Why AI Capability Is a Craft,
Not a Curriculum

Paul Gibbons | Adaptive Adoption™

paulgibbonsadvisory.com | June 2026

PILLAR 1 OF 7

DISCUSSION DOCUMENT – BETA

Rapid, emergent,
non-linear change.

1. The Craft Argument

Leadership is more like carpentry than calculus. So is AI.

You do not learn carpentry from a textbook. You learn it by cutting joints, ruining wood, adjusting, and cutting again. The knowledge lives in your hands and your judgment, not in your recall of terminology. AI capability works the same way. The dominant organizational response — design a curriculum, deliver modules, certify completion — produces people who can talk about AI but cannot build with it. That is the definition of a failed pedagogy.

Dreyfus and Dreyfus mapped this decades ago. Their five-stage model of skill acquisition — novice, advanced beginner, competent, proficient, expert — makes one thing devastatingly clear: formal instruction can get you from novice to advanced beginner, and no further (Dreyfus & Dreyfus, 1986). The transition from advanced beginner to competent practitioner requires *situated practice* — repeated engagement with real problems in context, under conditions of genuine uncertainty. Proficiency and expertise arrive only through years of pattern recognition built from thousands of concrete cases. No shortcut exists. No module substitutes.

This is not an argument against structured learning. It is an argument against the fantasy that structured learning alone produces capability. A six-week AI literacy programme is to AI mastery what a weekend woodworking course is to cabinetmaking: a reasonable start that most organisations mistake for the destination.

The word matters. "Literacy" implies a minimum standard — you can read, you can write, you are functionally adequate. "Fluency" is just faster literacy spelled with an F. Neither term captures what organisations actually need: people who know their tools intimately, understand their limits, apply them to solve real problems, and produce quality work through skill rather than luck. That is craft. A craftsman's knowledge is procedural, contextual, and built through doing. It cannot be transmitted through slides.

The practical consequence is architectural, not educational. You do not produce AI craftspeople by improving the curriculum. You produce them by creating the conditions under which craft develops: daily practice, sandbox environments where failure is safe, peer communities where knowledge circulates, and progressive challenges that push practitioners from one level to the next. The unit of learning is not the training course. It is the community of practice.



AI capability is not transmitted by curriculum alone. It develops through tools, materials, feedback, and repeated work on real problems.

2. Atrophy vs. Acceleration: The Bifurcation Nobody Is Measuring

Something is happening beneath the surface of AI adoption that most organisations are not tracking. Three distinct groups are forming, and the gap between them is widening faster than any corporate training programme can close.

The first group is atrophying. These are people who offload thinking to AI — not as co-creation, but as cognitive substitution. They paste a question, accept the answer, and move on. The pattern is seductive: AI outputs are fluent, confident, and fast. But fluency is not accuracy, and speed is not understanding. Over time, the muscles you do not use weaken. Cognitive offloading produces cognitive decline.

The evidence is accumulating. A meta-analysis presented at CHI 2025, synthesising seventeen studies on generative AI in higher education, found that AI tools produced large overall learning gains — but the effect was sharply attenuated for higher-order cognitive skills. Analysing, evaluating, and creating showed smaller gains precisely because students outsourced those

operations to the model rather than performing them (Qu et al., 2025). The finding is consistent: where AI replaces cognitive effort rather than scaffolding it, learners gain speed but lose depth.

Ferdman (2025) goes further, arguing in *AI & Society* that AI deskilling is not an individual failing but a structural problem. The concept of "capacity-hostile environments" names a condition in which AI mediation systematically impedes the cultivation of human capacities — not because users are lazy, but because the environment incentivises delegation over development. When the path of least resistance is to let the model think for you, most people will take it. That is not a character flaw. It is a design flaw.

A comprehensive review in *Computers in Human Behavior Reports* (2026) catalogues the hazards with clinical precision: endoscopists exposed to AI-assisted colonoscopy for twelve weeks showed a 21% decline in adenoma detection rates when the AI was removed. Airline pilots whose manual flying was limited to minutes per flight exhibited degraded situational awareness, with up to 44% failing basic cognitive tasks in simulators (Gerlich, 2026). The pattern is consistent across domains. Skills that are not exercised decay, and AI makes it extraordinarily easy to stop exercising them.

The second group is accelerating. These are people who co-create knowledge with AI — who treat every session as a learning opportunity, not just a task-completion event. For this group, capability compounds. Each interaction makes the next more productive because the practitioner is learning the grain of the tool: where it excels, where it breaks, which framings yield insight and which yield slop. They are not delegating cognition. They are augmenting it.

The Edo State Nigeria randomised controlled trial, conducted by the World Bank in mid-2024, is perhaps the cleanest demonstration. Eight hundred secondary school students used GPT-4 and Copilot in a structured six-week programme with teacher facilitation. The result: 0.31 standard deviations of learning gain — roughly equivalent to two additional years of schooling. Female students closed the gender gap entirely. The critical variable was not the tool. It was the structured pedagogy wrapped around the tool: teacher-introduced topics, carefully designed prompts, progressive challenge (World Bank, 2024). Access alone does not produce acceleration. Structured practice does.

The third distinction is where the premium lies. The first five dimensions of AI mastery (outlined below) are necessary. They produce competent practitioners. But the premium capability in an AI-saturated world is not competent use. It is *creative recombination* and *orchestration* — the ability to create something genuinely new with AI, and to coordinate human-AI ensembles that produce outcomes neither could achieve alone. These are not add-ons. They are the entire point. Organisations that train only for competence will produce users who never become architects.



The same tool can produce cognitive offloading or cognitive augmentation. The difference is not access; it is the discipline of use.

3. The Research Gap

Here is what no study Paul has read actually measures: *how* AI is used.

The research base treats AI use as binary — used it or did not use it. This made sense in 2023 when the question was whether people could benefit from AI at all. It makes no sense now. The difference between "summarise this document" and "co-build a multi-step analytical framework" is the difference between photocopying and authorship. Both involve a machine. They are not the same activity.

Dell'Acqua et al. demonstrated the "jagged technological frontier" in their landmark experiment with 758 Boston Consulting Group consultants — now published in *Organization Science* (2026). The finding was important: AI boosted performance on some tasks and degraded it on others, even within the same workflow and at seemingly similar difficulty levels. But even this sophisticated study treated AI access as present or absent. It did not measure the *quality* of interaction — whether consultants were using the tool as a cognitive partner or as a replacement for thought. The treatment was binary. Reality is a spectrum.

This matters because organisations measuring "adoption" by login counts, prompt volume, or licence utilisation are measuring the wrong thing entirely. A person who generates fifty low-quality prompts a day and accepts every output uncritically is counted as a power user. A person who spends forty minutes in a single, carefully structured session – iterating, challenging, verifying, and building genuine understanding – may appear less "engaged" by the dashboard metrics that most organisations track.

The field needs research that differentiates between cognitive delegation and cognitive augmentation. It needs longitudinal studies that track not just whether people use AI but how their use evolves over time. It needs measures of prompting sophistication, verification behaviour, and the degree to which practitioners integrate AI outputs with their own judgment rather than substituting one for the other. Until that research exists, most claims about AI's impact on productivity, creativity, or learning are confounded by the unmeasured variable of use quality.

Using AI in January 2023 is not the same as using AI in June 2026. The tools have changed. The expectations should have changed with them. The "have it do your homework" phase is over.

4. The Seven Dimensions

The framework below describes what a complete AI practitioner needs. Seven dimensions of breadth, four levels of depth.

1. **Civics** – Understanding AI's role in society: regulation, democratic participation, educational impact, labour displacement. You cannot make informed decisions about AI in your community, your children's school, or your organisation without grounding in what the evidence actually says. This is not optional civic seasoning. It is the foundation for every other dimension.
2. **Metacognition and Critical Thinking** – Knowing when to trust AI output and when to override it. This requires calibrated confidence, awareness of your own cognitive biases and how they compound with AI biases, systematic verification skills, and the ability to distinguish between model confidence and actual accuracy. As the marginal cost of information drops toward zero, human judgment becomes the scarce resource.
3. **Applied AI Ethics** – Practical ethical reasoning in your own work, not abstract philosophy and not compliance theatre. What does "fair use" mean for what you are building? When does co-creation become cognitive offloading? How do you verify outputs before publishing them? The distinction is *applied* ethics – doing the work, not quoting Kant.
4. **AI Concepts and Mechanics** – Executive-level literacy on how AI systems actually work: enough to make informed decisions about tool selection, architecture, and risk. You do not need a PhD in machine learning. You need to understand why different models behave differently, what architectural trade-offs mean for your use case, and where the genuine limitations lie.

5. **Practical Skills** — Prompt craft, context engineering, workflow integration, tool selection, and the orchestration of multi-tool pipelines. This is the mechanical dimension — the hand skills of the craft. It scales from basic prompting to full autonomous-agent deployment. Critically, it includes the skill of *delegation*: specifying objectives, guardrails, and what "good" looks like with a precision that most professionals have never been required to exercise. AI can produce mediocre output very fast, and human busyness means that mediocre-at-speed often ships unchallenged. The antidote is disciplined briefing — naming constraints, providing exemplars, and defining acceptance criteria before the model generates a single token.
6. **Creativity and Innovation** — Using AI as a creative amplifier rather than a creative substitute. The shape of human-AI creative work is a double diamond: the human premium sits at *both* ends of the process. At the front end, imagination and divergence — what matters, what is possible, what has never been tried. Many practitioners report that their "creative dam has burst" because AI makes more achievable than ever before; the constraint is no longer technical capacity but the quality of the questions you ask. At the back end, convergence, judgment, taste, and quality control — knowing when to stop, filtering signal from noise, exercising the aesthetic discrimination that separates craft from content. AI handles the middle well: exploration, iteration, pattern recognition, rapid prototyping. The human premium is at the edges. Organisations that fail to develop both ends — not just ideation but critical curation — will produce volume without value.
7. **Collaboration and Orchestration** — Designing and operating human-AI ensembles. This is where social complexity meets technical complexity: defining decision boundaries, building escalation protocols, governing AI agents that operate at scale, and managing the organisational change that accompanies any serious AI integration. Multi-agent workflows do not manage themselves. One underappreciated consequence: bad human delegation skills are *worse* in the AI world, not better. With human colleagues, shared history and institutional context compensate for vague briefs — "write me a report" can succeed because the colleague knows the audience, the standards, the politics. AI lacks that shared understanding entirely. Delegation therefore becomes a teachable, measurable skill rather than an afterthought — and its absence becomes visible in ways that organisations can no longer ignore.

The first five dimensions are necessary. They produce competent users. Dimensions six and seven are where the premium lives. Organisations that stop at five will produce people who can use AI but never architect with it.

The four levels describe depth of capability. Level 0 (Atrophy/Avoidance) describes knowing AI exists but not engaging, or engaging only through cognitive offloading. Level 1 (Assistant) is using AI for existing tasks, slightly faster — the 1.2x acceleration ceiling. Level 2 (Workflow Redesign) is restructuring work entirely, doing things that were previously impossible, achieving 5–10x capability expansion. Level 3 (Architect) is designing systems where AI operates as an autonomous agent within human-defined guardrails — competing on fundamentally different models.

The jump from Level 1 to Level 2 is the one most organisations never design for. People plateau at Level 1 and call it success. The job of a craft-development programme is to make that jump visible and navigable.

5. What Organisations Should Do Differently

Stop building curricula. Start building environments.



Practice architecture for the agent era: real work, feedback loops, community, and human judgement wired into AI-supported workflows.

The evidence from every study cited above converges on the same conclusion: access to AI does not produce capability. Structured practice does. The Edo State RCT did not hand students a laptop and walk away. It provided teacher-facilitated sessions, curriculum-aligned prompts, and progressive challenge over six weeks. The CHI meta-analysis found that guided AI use produced effect sizes 0.83 standard deviations larger than unguided use. The variable that matters is not the tool. It is the practice architecture around the tool.

Four design principles follow.

First, daily practice beats periodic training. AI capability is perishable. A quarterly workshop decays within weeks. The organisations seeing real capability gains build AI into the daily rhythm of work: "100 prompts" challenges that force practitioners to explore the boundaries of their tools, AI Fridays where teams share what they built that week, sandbox environments where experimentation is the expectation rather than the exception. Practice must be constant, safe, and social.

Second, craft coaches beat IT trainers. The person teaching AI craft should be someone who builds with it daily, not someone who was handed a curriculum last month. This is the reverse-mentoring model: digital-native practitioners who are building at Level 2 or 3 coach executives and mid-career professionals who are stuck at Level 1. Expertise flows in the direction of practice, not seniority. The traditional L&D model – senior people design training for junior people – inverts the actual knowledge gradient in AI.

Third, treat delegation as a trainable competency. The ability to brief AI well – specifying what good looks like, providing exemplars, naming constraints and failure modes – is not an innate talent. It is a skill, and organisations that teach it systematically will outperform those that treat prompting as typing. Delegation quality is measurable, coachable, and one of the highest-leverage interventions available.

Fourth, measure capability, not completion. If your adoption dashboard tracks login frequency, licence utilisation, or course completion rates, you are measuring inputs, not outcomes. The diagnostic question is not "how many people used AI this month" but "how many people can do something today that they could not do three months ago." Mastery-level self-assessment, peer review of AI-augmented work product, and portfolio-based evaluation are harder to administer than a login counter. They are also the only metrics that tell you whether anything is actually changing.

The prescription is environmental, not educational. You do not produce craftspeople by lecturing them about craft. You produce them by building workshops where they practise, communities where they share, and progressive challenges that push them past the plateaus where most people stop.

References

Dell'Acqua, F., McFowland, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F. & Lakhani, K.R. (2026). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*. <https://doi.org/10.1287/orsc.2025.21838>

Dreyfus, H.L. & Dreyfus, S.E. (1986). *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. Free Press.

Ferdman, A. (2025). AI deskilling is a structural problem. *AI & Society*. <https://doi.org/10.1007/s00146-025-02686-z>

Gerlich, M. (2026). AI-overdependence and human cognitive decline: Hazards, evidence, and mitigation strategies. *Computers in Human Behavior Reports*. <https://www.sciencedirect.com/science/article/pii/S2451958826001764>

Qu, X. et al. (2025). Generative AI Tools in Higher Education: A Meta-Analysis of Cognitive Impact. *CHI EA '25: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3706599.3719841>

World Bank (2024). From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria. Policy Research Working Paper 11125.

Paul Gibbons is an AI adoption advisor, keynote speaker, and author of eight books on leadership and organizational change. paulgibbonsadvisory.com