Journal of Expertise
2026. Vol. 9(2)

ISSN 2573-2773

The Fate of Expertise in the Age of AI

B. Scot Rousse[1] and Massimo Scapini[2]

[1]Department of Philosophy, University of California, Berkeley, USA
[2]Anchorsite AI, USA

Correspondence: B. Scot Rousse, bsrousse@berkeley.edu

Abstract

Generative AI is putting pressure on human expertise. For now, the pressure falls less directly on the embodied skills of craftspeople than on the so-called cognitive knowledge work carried out with the assistance of computers. AI now enables novices to produce expert-level outputs while bypassing conventional skill development. In many cases, practitioners must supervise work they cannot themselves execute while maintaining responsibility for outcomes they did not directly produce. Drawing on the Dreyfus Skill Model and Winograd and Flores's framework for human-computer interaction, we provide an interpretation of the metamorphosis that expertise is undergoing today. We sketch three stages through which practitioners develop a new form of orchestration expertise involved in responsible human-AI collaborations: awakening, calibration, and reimagination. Our central claim is that expertise is undergoing a fundamental, but unstable evolution, transforming into an ability to orchestrate machine capabilities while providing the human intuition and judgment needed to guide, assess, and take responsibility for work in ways no algorithm can automate. We conclude by arguing that this new expertise remains institutionally fragile, since the same AI systems that make orchestration necessary may also erode the entry-level roles and apprenticeships through which the judgment needed to orchestrate them is formed.

Keywords

orchestration expertise, generative AI, Dreyfus Skill Model, human-AI collaboration, AI mediated skill

Introduction

We now live in a world where AI systems predict protein structures, draft legal documents, generate software at massive scale, and assist physicians in diagnosing disease (Abramson et al., 2024; Chen et al., 2021; Dehghani et al., 2025; Everett et al., 2025). Although it is still too early to discern the full scope, some early analyses suggest broad exposure of white-collar work to AI disruption, especially at the entry level (Brynjolfsson et al., 2025; Simon, 2025). Even if this preliminary economic data remains somewhat cloudy, it is clear that a major reconfiguration of the nature of expertise is already underway. Specifically, this reconfiguration is taking place in the roles most susceptible to AI mediation, the so-called cognitive or knowledge work done with the assistance of a computer. We are seeing the rise of the expert as AI orchestrator. This is someone whose role involves discerning and validating quality in work they may not themselves be able to execute, and maintaining responsibility for outcomes produced through human-AI collaboration. At the same time, the viability of orchestration expertise is not guaranteed. Deployment of AI so far appears to erode entry-level pathways through which the judgment needed to guide it has been historically formed.

In this article, we draw upon the Dreyfus Skill Model (Dreyfus & Dreyfus, 1988; Rousse & Dreyfus, 2021) and the approach to human-computer interaction outlined by Winograd and Flores (1986) in *Understanding Computers and*

Cognition to make sense of the changes in cognitive-work expertise unfolding today. Inspired by and complementary to the stages of the Dreyfus Skill Model, we describe three stages in the emergence of what we will call *orchestration* expertise: awakening, calibration, and reimagination. These stages trace a movement from the initial disclosure of AI's disruptive possibilities, through the careful calibration of human judgment with machine capabilities, to the reimagination of whole workflows and professional practices. With the emergence of generative AI, expertise is undergoing a fundamental evolution, transforming into an ability to orchestrate machine capabilities that go beyond one's own, while supplying the judgment and answerability that are so important in human practices and that, as we will argue, should not be turned over to the machines.

We first contextualize the transformations in today's expertise within the Dreyfus Skill Model. We explore how generative AI unsettles the Dreyfusian developmental picture by allowing novices and non-specialists to produce outputs that once required long apprenticeship. Next, through following the case of "Maya," an accountant learning to build AI-assisted business-intelligence systems, we trace these three stages as this new expertise takes shape in a single career. Finally, we turn to the more vexing conundrum this transformation raises. If AI is eroding formative pathways through which expert intuition has traditionally been formed, how do we develop the next generations of experts capable of orchestrating AI systems? Our answer is exploratory. Apprenticeship has to be rebuilt — partly with AI, through what we call AI-calibrated apprenticeship, in which AI is folded into a novice's formation, under the guidance of human mentors, so as to build experience-based judgment rather than shortcut its development, and partly against the grain of AI, through the cultivation of care, responsibility, and communicative competence.

The Structure of Expertise

The Dreyfus Skill Model emphasizes how learners can progressively let go of following explicit rules and procedures as they gain familiarity and skill in a domain. Consider the difference between a novice nurse and an experienced one. The novice typically follows procedural checklists step by step, attending to features that can be specified in advance: vital-sign thresholds, scheduled medication times, charted symptoms. The experienced nurse responds to shifts in a patient's breathing, color, or affect, noticing that someone is taking a turn for the worse, adjusting care in real time, often before the monitors register the change. Patricia Benner's classic study of clinical nursing portrays this kind of transition in detail (Benner, 1984). Stuart Dreyfus and Hubert Dreyfus mapped this developmental journey towards expert intuition, which Benner applies in her studies of nursing expertise (Dreyfus & Dreyfus, 1988; Rousse & Dreyfus, 2021; Benner, Tanner, & Chesla, 2009). A wider literature on expertise in domains such as teaching, firefighting, chess, and music has provided important additional evidence for and refinements of the Dreyfus Skill Model's emphasis on direct, pre-reflective responsiveness as opposed to explicit deliberation and application of rules (Klein, 1998; Sudnow, 2001; Sutton, 2007; Christensen, Sutton, & McIlwain, 2016; Mangiante, Peno, & Northup, 2021; Rousse, 2019).

Now, the Dreyfus account is one notable model of expertise among others, and it has been challenged on multiple fronts: by philosophers who argue that conscious, conceptual thought is more pervasive in expert action than the model allows (McDowell, 2007; Montero, 2016), by psychologists within the deliberate-practice tradition who argue that expertise in "elite" skills does not consist in the "smooth, skillful coping" which the Dreyfus model sees in "everyday action" (Dreyfus, 2014), and instead involves effortful training designed to improve specific aspects of performance under guided feedback (Ericsson et al., 1993; Ericsson, 2006). Finally, the Dreyfus Model has been forcefully and directly challenged by the expertise researcher Fernand Gobet, who argues that expert performance depends on discrete knowledge structures such as chunks and templates that support rapid pattern recognition and guide selective look-ahead search (Gobet & Chassy, 2009; Gobet, 2015). Indeed, Gobet (2018a) explicitly argued in the pages of this journal that the Dreyfus account implausibly, even fatally, underestimates the role of analysis and search in expert cognition.

However, the Dreyfus Skill Model has resources to answer these objections. For example, Dreyfus and Rousse (2018) present detailed counterarguments to Gobet's account, prompting a reply from Gobet (2018b). Rousse and Dreyfus (2021) then respond further, answering Gobet and addressing the broader set of critics mentioned above. In the case of the deliberate-practice tradition, in particular, the disagreement is narrower than it first appears: The Dreyfus Skill Model holds that experts, in order to keep learning, deliberately step back into effortful engagement, so that Ericsson's deliberate practice can be incorporated as continuous with the Dreyfus picture rather than treated as a rival to it. (This has the correlative result of undermining Ericsson's distinction between elite and everyday skills.) The Dreyfus Skill Model rightly remains an influential and well-grounded account of skilled performance (see Mangiante et al., 2021), and even though Dreyfus and Dreyfus (1988) were critical of attempts at AI automation of human skilled work (e.g., the "expert systems" of the 1980s), we will argue here that their skill model provides an illuminating framework for analyzing and navigating how the automation enabled by today's generative AI is reshaping the nature of expertise.

The Dreyfus Skill Model proposes that learners typically progress through five stages, with a sixth stage of mastery available to those highly motivated to continue beyond conventional expertise. *Novices* follow explicit rules applied to objective features they can recognize without previous experience. *Advanced beginners* begin to orient to situational patterns gleaned through repeated exposure: the sound of an engine being pushed too hard, the smell of oil beginning to burn. *Competent* performers develop the ability to choose a plan or perspective that organizes the situation's complexity, rendering some elements relevant and others ignorable, and heightening emotional involvement in outcomes. *Proficient* performers directly grasp situations in light of a perspective that intuitively dawns on or occurs to them, revealing what is relevant and guiding their approach to the situation, though they still deliberate about execution. For *experts*, action in familiar and intuitively grasped situations flows directly, without deliberation. As the Dreyfuses observed, "when things are proceeding normally, experts don't solve problems and don't make decisions; they do what normally works" (Dreyfus & Dreyfus, 1988, pp. 30–31). At the sixth stage, *masters* proactively adapt to changes in the skill domain, including those wrought through new technologies. Masters deliberately experiment to expand their repertoire of intuitive responses, thereby pushing the boundaries of conventional expertise and sometimes discovering new possibilities that reshape entire domains (Rousse & Dreyfus, 2021).

The Dreyfus Skill Model focuses on an individual practitioner's developmental arc. But most professional knowledge work — like all cooperative human activity — also requires skills for coordination among multiple participants (see Rousse, 2026b). Drawing on John Searle's (1979) taxonomy of speech acts, Winograd and Flores (1986) argue that organizations (both informal and formal) are constituted as networks of recurrent conversations, networks of offers, requests, declarations, and promises through which people coordinate their actions by making and following through on commitments. On this expanded view, expertise requires and draws upon the capacity to participate competently in the conversational networks that enable them to participate effectively in their organization and practice domain. This "communicative competence" (Winograd & Flores, 1986, p. 163; Flores, 2012) includes skills for listening for what is needed, negotiating understanding, delivering upon commitments, and bearing responsibility (making amends, restoring trust) when things go wrong. As we will see, this second dimension of expertise is relevant because tomorrow's professional AI orchestrators must focus not only on individual performance but also on vouching for results and answering for breakdowns or lapses in quality.

Skill domains have always evolved as their characteristic technologies have changed. For example, when tennis rackets shifted from wood to graphite and increased in head size, players had to adapt their technique (Rousse & Dreyfus, 2021). A similar change happened when architectural drafting moved from paper to CAD software, when legal research moved from law libraries to digital databases. Sometimes skill domains also get completely disrupted and marginalized with the emergence of a new technology, as happened, e.g., when cinema production and exhibition switched from chemical photography to digital.

But generative AI brings something different. In domain after domain, people with neither formal training nor years of apprenticeship can now use AI to create work that previously required the developed powers of experienced practitioners. These tools look like they may become the kind of radical innovation that, in the words of Winograd and Flores (1986, p. 6), opens up "whole new domains of possibilities in the network of human interactions." Along the way, these technologies raise the question: Is the hard-won identity of the traditional expert a relic of a bygone era? And is the orchestration expertise we describe a stable formation or a transitional phase to a more fully automated and technologically controlled era of work, one where human judgment plays less and less of a role? The answers to these questions depend partly on what we as a society do with these new technologies, how we integrate them into our practices.

Before proceeding, it will help to gather a few key definitions. These interrelated terms are integral to or follow from the Dreyfus Skill Model and are important for our account of orchestration expertise; some have already appeared in the preceding section.

Intuition. The Dreyfus Skill Model uses "intuition" as a technical term. To have intuition in a domain, or to be intuitively familiar with it, is to "know your way around" in it. That means you have had enough experience with its recurrent situations such that you have a feel for the kinds of things that tend to happen, so that a typical situation "makes sense" to you; you see what to do, rather than feeling overwhelmed with a morass of competing details. In the terms of the Dreyfus Model, this is to say that a *perspective* on the situation non-deliberatively dawns on you. A perspective (sometimes also called a "plan") is a way of relating to a situation that brings together an orientation toward a goal, a differential salience among its elements (some standing out as relevant, others receding into the background), a provisional sequencing of the steps of the task, and the performer's emotional investment in the outcome. For example, a competent cook preparing several dishes for a meal does not simply follow separate recipes in serial order but chooses a perspective that organizes the preparation of the whole meal, such as having the cold salad ready before the hot pasta. This perspective gives the cook a goal, making some elements of the kitchen salient, e.g., the vegetables, pots, burners, and salad spinners, while letting unrelated ingredients, unused tools, decorative choices, and other possible tasks recede into the background. The perspective also sequences the sub-tasks so that the salad remains crisp while the pasta arrives hot. Intuition has a diagnostic edge — an immediate, pre-reflective sense that something has started to go wrong, or that a once-familiar situation has turned novel and now calls for deliberate attention (Rousse & Dreyfus, 2021).

Judgment. In a word, with the term "judgment" we mean to capture "what we get at (or should be getting at) when we say that someone 'has good judgment': a form of thinking that is reliable, just, and committed" (Smith, 2019, p. xvi). Judgment thus involves the ability to see what to do in a situation, what is good, worthwhile, or important to do. It involves being not only able to tell how things are going in a situation — how they might be going wrong, how they might be improved — but also to grasp why. It is a form of reflective thought that attends to the specificities of a particular case rather than only to the general pattern it instantiates (Benner, Tanner, & Chesla, 2009). It also involves a form of refined awareness of, and respect for, the standards of the given domain, including when a situation calls for breaching or revising the standards (Haugeland, 1998). Coaches and teachers usually have the talent to articulate their judgment in language, answering questions and engaging in conversation about it, even while appreciating that any explanation can only approximate, never exhaust, the full background of intuition from which their judgment arises (Taylor, 1995).

Discernment. Discernment is the evaluative dimension of judgment: the capacity to tell what is worth doing; to make fine distinctions of quality and taste, to tell good work from passable and excellent from good, and to distinguish what genuinely fits a situation or rises to an occasion from what is merely procedurally adequate.

In sum, intuition carries a practitioner when things proceed normally in familiar situations, and judgment and discernment enable one to navigate situations when they do not, e.g., when a situation is novel, there is some hesitation about what to do, stakes are high, or one is evaluating someone or *something* else's output. Evaluating AI output is exactly such a case. Today's new orchestration expertise is reoriented from direct performance in intuitive flow toward the judgment and discernment required to guide, assess, and correct work one did not personally execute but delegated to an AI system. Whatever the AI produces, it is the human practitioner who must guide the process, verify and vouch for the result, and answer for what goes wrong. As Andrej Karpathy, one of the founders of OpenAI and a major voice in today's AI-induced transformations in skill and work, recently put it with respect to today's new mode of software engineering: "You are still responsible for your software, just as before . . . You still have to be in charge of aesthetics, judgment, taste, and oversight" (Karpathy, 2026). With this stage-setting in place, we can now proceed to the details of our positive account and illustrative case study.

Expertise in Transition

If the Dreyfus skill model describes a ladder climbed rung by rung over years of involved practice, generative AI offers something more like an elevator that can take you directly to higher floors, bypassing the traditional developmental journey. The question of whether you know what to do once you step off the elevator points to the need for a new form of skill. The shifts taking place here become especially palpable when seen in unexpected contexts. For example, one of us watched his eight-year-old daughter create a complete song in a single afternoon, using ChatGPT to write lyrics and Suno to compose and produce music. Within hours, her creation was streaming on Spotify alongside professional recordings. A friend who listened asked: "Nice song, but did she actually learn anything?"

Of course, performance gains do not always involve learning. A marathoner shod in carbon-plated racing shoes runs faster, but no one would say the shoes have taught her anything. The technology lifts her output without enriching her skill. So, one might wonder whether the daughter's song is similarly a gift of the tools, with no actual learning behind it. We don't think so. While she certainly did not develop musical expertise in the traditional sense, she did equip herself with a skill just coming into focus. She learned to prompt and guide the tools, and to discern when the lyrics captured or failed to capture what she wanted to express. While not rising to the level of fluency of an expert composer, this is nevertheless an emerging domain of skill that involves its own demands of intuition and judgment.

Another anonymized case based on real consulting experience helps bring home the stakes of the transformations underway. During a demonstration of AI capabilities to a construction project manager with fifteen years of experience, an LLM analyzed a complex coordination issue and drafted a response in seconds. The manager's reaction: "That's not my voice. People will know I didn't write this."

He did not want to use the machine's output. His professional reputation had been built on being the person who knows what to say and how to say it in these situations. The very identity of the traditional expert is seemingly under threat. Many people are resting content with the attitude that "I don't need AI; I'm an expert." But what if the more valuable expertise now involves discerning when AI can be relied upon?

Still, AI skepticism is well-grounded. A recent report by Challapally, Pease, Raskar, and Chari (2025) on enterprise AI implementations found that ninety-five percent flopped, failing to produce measurable improvements in business operations. Their account suggests that these failures often stem from treating AI as a form of "plug-in automation," and thus as a path to eliminating workers, rather than as a new domain requiring new forms of human expertise. The difficulties involved in expertly incorporating AI into existing roles and workflows are corroborated and ramified by a recent study of "workslop" (Niederhoffer et al., 2025). The study finds that when organizations push rapid AI adoption without cultivating practices of meaningful human oversight, workers begin to be inundated with low-quality automated drafts that require substantial effort to correct.

These developments signal the need for a new approach to expertise. Making responsible use of generative AI is itself a novel form of expertise. It is an expertise focused on guiding, assessing, verifying, and taking responsibility for work produced in collaboration with AI systems. These capabilities rest on domain intuition that still requires active cultivation. Realizing this means that organizations should stop treating AI as plug-in automation and redesign processes around how best to guide it. This recognition is already taking hold in some corners of industry. Again, Karpathy drives this point home with respect to software engineering: "Even if [AI] agents do more of the work, the human still needs understanding to direct them. You need to know what is worth building, what question matters, what result is suspicious, and what tradeoff is acceptable" (Karpathy, 2026). Julie Sweet expands this point to the organizational level. Regarding Accenture's reorganization around AI capabilities, Sweet observes that "it's not the technology that is the biggest barrier, it is actually being able to get the mindset re-organized around how best to use it, the ability to do the change management, the process reinvention" (Lichtenberg, 2025).

Portrait of a Transformation: Maya's Journey

To make this transformation more concrete, we offer the story of Maya, an accountant at a restaurant chain. Maya is an anonymized case drawn from an actual engagement one of us led. She originally dismissed AI as "cheating." Today she credits it as essential to doing work she could not previously have imagined doing alone. Her evolution began with a startling demand from leadership: within three months, create a real-time revenue analysis system in Power BI, a software platform for interactive data visualization. She had never used Power BI, and her fifteen years of restaurant accounting had not prepared her for data architecture or programming.

The assignment seemed impossible. "It would have taken me two years of courses to even become a beginner in Power BI," she explained. She knew this from experience. Having learned Excel through a series of certifications, she understood the limitations of that path. Certifications alone don't create skill: one needs further practice and guidance. But consultations from Power BI experts are prohibitively expensive. So, she skipped the traditional certification route entirely and decided to turn to ChatGPT. Her early attempts failed spectacularly. "The reports never matched what I wanted. I didn't know how to ask the right questions." She would paste ChatGPT's code directly, watching nothing work as promised. Frustration nearly drove her to abandon the technology entirely.

Her breakthrough came when she changed how she used the technology. She realized she should use AI to learn. She could treat ChatGPT as a dependable tutor for twenty dollars a month. But using ChatGPT to learn by practicing only works if you're moving in the right direction, avoiding hallucination loops. This discernment is where her domain expertise proved pivotal: she could validate everything the AI produced against what she already knew: her Excel reports and the business's accounting system. Today she generates complex Power BI code without understanding programming syntax in the way a computer science graduate would. Her domain familiarity allows her to verify outputs against known benchmarks. "Sometimes AI gives me results for the whole fiscal year when I asked for specific days. That is when my expertise matters: I catch those errors." In addition, her new capability here is connected with learning to ask ChatGPT to teach her rather than just to do something for her. She learned to ask, "What are the steps I need to take?" rather than simply "Do this."

But Maya did not succeed alone. When repeatedly hitting walls with circular AI answers and half-built dashboards that wouldn't connect to live data, she discovered a WhatsApp group of professionals navigating the same challenge. "ChatGPT can write code," she explained, "but it can't tell you that your company's specific database structure needs a different approach. Only someone who's been there can do that." When she couldn't figure out how to structure her data model, a controller in Mexico City walked her through his approach. In these groups, AI-enabled tinkerers support each other without traditional master-apprentice relationships or structured curriculums. Organizations have not yet widely created these support structures. Practitioners themselves are improvising them because the available tools outstrip the available training. In the process, they are creating communities of practitioners who are learning together how to make AI work with their previous domain familiarity, building the judgment required for responsible AI orchestration.

Maya's story illustrates the first two stages in the emergence of orchestration expertise: awakening and calibration. Her initial resistance to AI, followed by the recognition that her existing accounting expertise no longer sufficed for the task before her, marks the moment of awakening. Her subsequent practice with ChatGPT, her repeated testing of AI-generated outputs against Excel reports and accounting records, and her participation in a peer community of AI tinkerers mark the work of calibration. The third stage, reimagination, appears less fully in Maya's individual trajectory. As we will show, reimagination becomes visible when calibrated practitioners and organizations begin asking whether the inherited artifacts, workflows, and roles of a domain should keep their present shape at all.

The three stages we propose, awakening, calibration, and reimagination, presuppose, elaborate, and extend the Dreyfus picture into a new technological situation. They track not how a novice develops domain expertise from scratch, but how a practitioner who already has some degree of domain familiarity and skill navigates the disruption and reconstitution of that skill in the presence of AI as a new partner. Each stage has a counterpart in the Dreyfus vocabulary. Awakening corresponds to a performer's recognition that she has encountered a sufficiently novel situation in which her familiar intuition is no longer reliable, leading her to adopt a deliberative approach. Calibration is the process through which practitioners move through *competence* to *expertise* in incorporating AI tools in their practice: targeted practice, asking for guidance, following procedures, taking responsibility for choices, being invested in outcomes, and gradually moving toward the ability to meaningfully orchestrate AI collaborators. Finally, reimagination echoes what the Dreyfus Model calls *mastery*: proactively experimenting to reveal new possibilities of performance rather than simply doing what normally works.

Stage 1: Awakening

Awakening begins with a visceral recognition that a new technology has shifted the ground beneath one's professional identity, and that one can no longer ignore it. It is less a stage of practice than the threshold of the transition, a reorientation toward the domain's altered reality rather than a span of active skill-building.

The Dreyfus Skill Model illuminates this moment of awakening. Expert intuition is built up, over years, from repeated encounters with the recurrent situations of a domain, and it works by recognizing those situations and responding to them without deliberation. So long as the situations a practitioner meets are familiar ones — so long, in the Dreyfus brothers' phrase, as things "proceed normally" (Dreyfus & Dreyfus, 1988, p. 30) — that intuition carries the practitioner without any need to deliberate. Generative AI breaks this condition: A new piece of equipment has transformed the typical situations of the domain, so that what a practitioner's intuition was formed to read is no longer what they encounter. Intuition loses its grip, and the practitioner, who once simply acted, must now stop and deliberate. That loss of footing is what Maya underwent, as we saw, when she was handed an assignment that her expertise in restaurant accounting had not prepared her to accomplish, though her knowledge of restaurant accounting would enable her to know if she succeeded with the assignment.

What awakening discloses, first, is that competence and expertise have been redefined. In the cognitive and knowledge work we are concerned with, being good at the job increasingly means being able to discern when AI-generated work meets the standards of the domain, even when one could not have produced that work oneself. It drives home, second, that resting content has itself become a risk: The conviction of an accomplished expert that he has no need of AI can become the harbinger of obsolescence rather than a mark of professional confidence. Finally, the moment of awakening hints that the practitioner will not meet the change alone. Expertise is becoming more distributed, drawing the practitioner into working relationships both with AI "partners" and "interns," and sometimes additionally with communities of fellow practitioners who help make sense of unfamiliar tools and work through breakdowns together. How those relationships are built is the work of the stage that follows.

Stage 2: Calibration

If awakening is recognition, calibration is the work that recognition sets in motion: the deliberate, effortful practice through which the disrupted practitioner builds the new competence of an AI orchestrator. Here calibration draws on deliberate practice (Ericsson et al., 1993), discussed earlier in connection with the Dreyfus Skill Model — practice that is structured, focused on specific weaknesses, and guided by feedback that shows where one falls short — and adapts it to a new object. The practitioner is no longer refining a single skill at the limit of her own ability; she is learning to work with a new kind of partner, at once coach, teacher, and unreliable colleague, whose capacities are extraordinary in some respects and erratic in others. The feedback of this deliberate practice, accordingly, runs in two directions: from the AI's outputs back to the practitioner, and from the practitioner's domain knowledge back into the prompts that draw better output from the AI.

Calibration of AI competence, so understood, is made up of several interlocking practices, and Maya's case, told above, shows them taking hold. First, the practitioner practices and experiments in low-stakes settings such as internal reports and prototypes, where errors can be caught before they reach an operational or financial decision. In this zone of safe tinkering the practitioner maps the powers and the limits of the tool without work or reputation at stake.

A second crucial AI calibration practice is establishing verification and responsibility rituals. The "jagged" nature (Dell'Acqua et al., 2026; Karpathy, 2026) of today's generative AI systems, that is, their aforementioned vexing combination of virtuosity and idiocy, and their tendency to confidently confabulate statements and figures, means it is imperative to hold their output to standards the practitioner already commands, as Maya did when she required her dashboards to match her Excel records. In calibrating to AI's capabilities, orchestrators in training must hold fast to the fact that they themselves are the final guardians of the quality of the AI output. The fact that it can be so easy to prompt an AI to generate multiple outputs can result in an emotional distance from that output, but this emotional distance must be kept closed.

A third calibration practice pertains to what the practitioner asks the AI to do. Many beginners with generative AI ask for answers and solutions, as they would an old-fashioned internet search. The person calibrating their ability to work with and orchestrate AI systems learns the power of asking "how to" questions instead. They ask, "What are the steps I need to take?" rather than, "Do this." Fourth, in calibrating AI competence, it is important to interact not only with AI systems and agents, but with other people who are interacting with AI systems and agents. The calibrating practitioner finds groups of fellow tinkerers who share both best practices and horror stories and can serve as AI adoption mentors. Recall Maya's WhatsApp group, and the controller in Mexico City who walked her through her data model. AI calibration should be approached as a social endeavor.

Gradually this practice yields a fluency at which one can comfortably delegate and direct the technology's capabilities toward ends one sets. At this point one has become an orchestrator. Karpathy describes a version of this figure and himself supplies the terminology of the *orchestrator*. In his account of frontier software engineering, what he calls "agentic engineering," the developer using a system such as Claude Code is no longer chiefly an author of code but an "orchestrator of [AI] agents," directing fallible AI "interns" that "fill in the blanks" while the engineer provides the direction they need (Karpathy, 2026).

We are using the term *orchestrator* in a broader sense to capture how any person using an advanced AI system can learn to orchestrate its capabilities, whether or not this involves AI coding agents. In our sense, Maya orchestrates ChatGPT to produce Power BI code, and the daughter in our story above orchestrated Suno and ChatGPT to create a song.

Calibration, as we have described here, is the path of a practitioner whose domain expertise predates their adoption of AI tools and has been unsettled by it. Those who initially enter a domain after generative AI has already begun to transform it undergo a different version of the stage. We return to this emerging form of AI-calibrated apprenticeship below.

Stage 3: Reimagination

If calibration is the rebuilding of competence, reimagination is what that competence, once rebuilt, makes possible. Becoming an orchestrator is not yet reimagining. The calibrated orchestrator still works within the existing shape of her job, doing the work her role already comprises, only now with AI's assistance and with the expanded reach it provides. Reimagination begins when the orchestrator asks whether the skill should keep its present shape at all.

Karpathy again furnishes a vivid and fresh case study here. The temptation, he warns, is to treat AI merely as "a speedup of what exists," a means to work faster within an inherited paradigm rather than to question it. The example he provides is the way he sees software coding now moving away from a focus on applications ("apps"). A standing application is a fixed artifact, assembled in advance to deliver a capability on demand. To make such a software capability available to people used to mean building an app to embody it, a dedicated piece of software written and maintained for that purpose. The example he gives is relatively trivial, but it makes a dramatic point. He describes creating an app that automatically adds pictures of dishes to a given menu (once you upload a picture), enabling diners to better visualize their options. But when generative AI can perform tasks and deliver capabilities directly, as he found Google's Gemini could do in this case, producing what is needed upon request, dedicated apps are no longer needed. Thus, he concludes, the "app shouldn't exist" (Karpathy, 2026).

A question for anyone experimenting today with generative AI is this: What analogs to Karpathy's described "paradigm shift" in computing might now be available in your own domain of practice? What artifacts, workflows, and constraints are becoming irrelevant? What are the contours of this reconfiguring horizon of possibility? Again, acknowledging that software engineering is a domain that is particularly susceptible to this kind of AI-induced reconfiguration, we believe the possibilities for reimagination are available far and wide and that further structured experimentation will be required to bring these possibilities out.

While individuals reimagine their own work so that they leave some activities to AI agents, organizations do the same for whole work processes. For example, Ruthanne Huising's studies of organizational process-mapping show that formal workflows are, in practice, dense webs of workarounds and informal negotiations (Huising, 2019). To layer AI onto a workflow like this without first remaking it would likely just replicate its existing waste at higher speed. Reimagination asks instead what the work would look like if it were designed afresh in light of what AI now affords. This is the reinvention Julie Sweet had in mind in observing that the binding constraint on AI adoption is the reorganizing of process and mindset, not the technology itself (Lichtenberg, 2025). Some organizations have begun to institutionalize this mode of reinvention. When the home-furnishings retailer IKEA introduced an AI assistant to field routine customer-service inquiries, the company simultaneously retrained roughly 8,500 call-center workers as remote interior-design advisers, recognizing that freed capacity could support more complex, consultative interactions. The AI assistant handled simpler queries, while workers were empowered to take on advisory work that called for design competence, relationship-building, contextual sensitivity, and judgment, building around them a remote design-advisory service that had not previously existed at such scale (Ingka Group, 2023; Reid, 2023). The routine load passed to the machine, while the work that called for taste, context, and judgment was entrusted to people. Arguably, AI systems could also be integrated into these design practices, with an experienced designer on hand to orchestrate them.

The Vanishing Bottom Rung

New roles created by reimagination still require the cultivation of domain experience and judgment to be responsibly implemented. If AI innovations are driven and incorporated in such a way as to displace more and more human practitioners, restricting or precluding the development and social transmission of domain intuition, where will the next generations of experts needed to orchestrate AI systems come from? Recent research indeed suggests that in AI-exposed knowledge work, early-career roles are contracting, driven by a slowdown in junior hiring (Hosseini Maasoum & Lichtinger, 2025). Corroborating these findings, research by Brynjolfsson et al. (2025) indicates that, so far, generative AI reduces employment most steeply at the entry level, while senior positions remain comparatively stable. The ultimate explanation of these trends is actively debated. If the above-described pattern holds, the bottom rungs of the ladder of experience and expertise are beginning to weaken and could break. To respond to these shifts, the early stages of skill and career development should also be reimagined. Concretely, organizations will need to create protected novice roles built around bounded but consequential assignments, routine human review, shared libraries of AI failure cases, peer calibration groups, and explicit verification rituals that teach novices not merely to use AI, but to become answerable for what AI helps them produce.

One element of the requisite reimagining is what we call *AI-calibrated apprenticeship*, a rearrangement of early learning in which use of AI becomes included in a practitioner's formation rather than a way of bypassing it. It is, in effect, the calibration stage of our above account adapted into a path for those who never held a pre-AI competence to recalibrate. The guiding principle here is that AI-mediated learning should be organized to foster experience-based judgment rather than to shortcut its development. The learner still needs repeated contact with the typical situations of the domain but now meets them partly through and with AI: comparing its outputs against known examples, receiving expert human feedback on why a given AI output is adequate or deficient, and practicing diagnostics of AI failures rather than merely accepting corrections.

A novice accountant need not perform every reconciliation by hand before working with AI.

They need to receive enough first-order training combined with human-mentor-guided exposure to generative AI's powers and idiosyncrasies so that they can learn to discern trustworthy output and attune to the sense of unsettlement that reveals discrepancies, exceptions, and confident-sounding AI errors. In this way, AI-calibrated apprentices can still develop the intuition needed to orchestrate, assess, and validate AI output. They would cultivate the foundations of intuitive familiarity and judgment: the capacity to recognize a typical situation, to tell what works from what looks wrong, and to sense what calls for further scrutiny.

Again, the "jagged" nature of today's systems requires that AI-led trainings always be conducted with regular guidance by and contact with a human mentor or peer community who can narrate breakdowns, share verification rituals, and supply situational perspective regarding how the AI systems are likely to fail and how they can be most effectively used. This calls for designated spaces in which novices can work on consequential but bounded problems, seeing far more cases than traditional apprenticeship could offer, and receiving feedback from both machines and people. But even this is not enough.

Responsibility and the Fate of Expertise

Today's LLM-based AI tools can multiply examples, simulate scenarios, stage practice, offer guidance, and expose a learner to multifarious failure modes. But AI systems cannot teach what tasks and projects are worth undertaking, or how to care about results, how to take responsibility for mistakes, or how to preserve and rebuild the trust that enables cooperative human endeavor. Here the developmental account we have drawn from the Dreyfus Skill Model meets the second, social dimension of expertise brought into focus by Winograd and Flores and reconstructed recently as part of the larger phenomenological tradition in the philosophy of work (Rousse, 2026b). An apprenticeship adequate to expertise in the age of AI should cultivate not only the intuition to recognize and validate good work, but this social and conversational dimension of expertise as well.

Cultivating a skill is usually not a private endeavor. To exercise a skill is to participate in a domain of practice in which the relevant skilled actions have their place (Wrathall, 2017). Further, to participate in a domain of practice is to be subject to pre-existing standards one has not oneself authored, e.g., standards for what count as a good performance, or a legitimate move in the relevant game (MacIntyre, 2007). A chess player cannot decide for herself what counts as a legal move or a strong position. The accountant dedicated to her practice holds the standards of sound accounting as binding, treating breaches as demanding honest repair rather than glossing them over or letting them slide (Haugeland, 1998).

Participating in a domain of practice also means participating in the networks of relationships and conversations that enable the people carrying on the practice to conduct its characteristic cooperative activities, both informal (e.g., meeting a friend or bandmate for practice) and formal (competing in an official tournament or competition). Coordinating these activities involves a structured "dance" of recurrent conversations: requests, offers, promises, declarations, and assertions (Winograd & Flores, 1986, p. 64; Flores, 2012). In this picture, communication is not merely a *transmission of information*; communication is the *coordination of commitments* generated by the relevant speech acts. For example, a promise is a commitment personally to make something happen such as creating a software app in order to take care of someone else's concern. Delivering on a promise requires the fulfillment and coordination of commitments by both promisor and promisee. An engineer's commitment to deliver a software application can only be fulfilled if the client has implemented the required computational infrastructure to properly host the app.

One benefit of this account of the conversational coordination required to exercise skills is that it gives definite content to all of the above talk about "responsibility." We have emphasized over and over again that AI-orchestrators must assume responsibility for the quality of the AI-generated output they usher into the world. When we say the orchestrators are responsible, we mean much more than that they played a causal role in generating the output by prompting the AI. We are pointing out that the delivery of the software amounts to a promise they are making to the one receiving the software: It is a promise that the software will take care of the concern it was designed to address and that it will work in the agreed upon way; it is, moreover, a promise to stand behind what was delivered, to facilitate support and repair if it does not work; it is a promise to engage in the requisite follow-up conversations to better understand the underlying animating concerns in order to improve delivery of subsequent iterations; and it is a promise to listen, make amends, and offer appropriate recompense in order to rebuild trust if the project fails through the engineer's fault (cf. Flores, 2012). Whether future AI systems will eventually be able to exhibit a compelling degree of normative competence in human relations is ultimately an empirical question (Rousse, 2026a), but until this social and conversational responsibility can also be simulated by a machine, all predictions about AI systems soon taking over all human work, or even all human knowledge work, are just speculation or hype.

However, no matter how AI capabilities evolve, responsibility of this kind cannot be completely delegated to any AI system, that is, unless we, for some reason, aspire to take human beings out of the conversational "loops" in which our very concerns are being interpreted and addressed. But if we tried to do so, we would be creating for ourselves an AI-powered intensification of the bureaucracy described by Hannah Arendt in the last century: "In a fully developed bureaucracy there is nobody left with whom one can argue, to whom one can present grievances" (Arendt, 1970, p. 81).

In a word, what participation in the above-described conversational networks asks of a practitioner is *care* (Rousse, in press). In the sense at stake here, care is not just a feeling of affection (Denning & Rousse, 2024); it becomes real in how well one tends the conversations, commitments, and relationships through which a skill is concretely carried out, preserved, evolved, and then passed on to the next generation. The fate of expertise in the age of AI and the viability of the expert as AI orchestrator, then, are inseparable from the fate of the skills through which human beings exhibit the conversational care and responsibility discussed here.

People can obviously be better or worse at listening to each other's concerns and taking care of these concerns in conversation (e.g., making and delivering on an offer or request in a timely way). But this is not a fixed given; as Winograd and Flores explicitly argue, these conversational skills can be cultivated, refined, or allowed to deteriorate. By extension, again as Winograd and Flores argued, computer and AI systems integrated into organizations can be better or worse both at facilitating the conversations that cooperative work requires and at supporting the development of people's own "communicative competence" and attunement to commitment (Winograd & Flores, 1986, p. 163). This attunement should be a central concern for the design and deployment of today's AI systems.

The AI orchestrator, then, is not simply a more efficient user of powerful tools. The orchestrator is a practitioner who has been formed in a domain deeply enough to know what matters, to recognize when a machine-generated result is adequate or deficient, and to stand behind the commitments that those results put into circulation. Orchestration expertise exists only where institutions preserve and redesign the practices through which people become answerable: setting the goals and purposes, mentoring, review, correction, repair, respect for shared standards, and participation in communities where the consequences of one's work matter to others. Such a figure cannot be produced by access to AI alone.

Ultimately, whether the expert as AI orchestrator proves a lasting figure or a transitional one is not a technical question. The orchestrator could turn out to be a mere placeholder on the way to a future in which more and more skilled, coordinated human activity is given over to AI, the ladder of expertise is left in ruins, and we become, in Stuart Russell's words, enfeebled "passengers in a cruise ship run by machines" (2019, p. 255). How things turn out will depend on whether we take a stand, in our society, in our organizations, and in our own lives, to preserve the impassioned, embodied, and relational practices through which skilled work, judgment, responsibility, and care are formed. That means continuing to mentor newcomers in entry-level roles redesigned around AI, continuing to help them learn, honor, and evolve the standards of our domains of practice even where AI performs much of the work. Human intuition, formed through care and experience, is also what lets us sense which tasks are worth undertaking, which projects deserve our commitment, and which possibilities should be allowed to matter in the first place. The fate of expertise remains in our hands.

Acknowledgements

We would like to thank Stuart Dreyfus, Fernando Flores, and Terry Winograd for many conversations that shaped our thinking and this essay. We are grateful to Guillermo Campitelli and Charles Spinosa for valuable detailed feedback and suggestions on earlier drafts. Especially given our subject matter, we would like to acknowledge that we orchestrated commercially available LLM tools in the drafting and review of this essay. We take full responsibility for all claims and content. Finally, Massimo is grateful to his daughter P. S., who served as the principal collaborator in the family AI experiments out of which much of his thinking here originally grew.

ORCID iDs

B. Scot Rousse
https://orcid.org/0000-0002-6840-9641

Massimo Scapini
https://orcid.org/0009-0004-8800-0722

References

Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., Bodenstein, S. W., Evans, D. A., Hung, C.-C., O'Neill, M., Reiman, D., Tunyasuvunakool, K., Wu, Z., Žemgulytė, A., Arvaniti, E., ... Jumper, J. M. (2024). Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature, 630*(8016), 493–500. https://doi.org/10.1038/s41586-024-07487-w

Arendt, H. (1970). *On violence*. Harcourt, Brace & World.

Benner, P. (1984). *From novice to expert: Excellence and power in clinical nursing practice*. Addison-Wesley.

Benner, P., Tanner, C. A., & Chesla, C. A. (2009). *Expertise in nursing practice: Caring, clinical judgment, and ethics* (2nd ed.). Springer.

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/

Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025, July). *The GenAI divide: State of AI in business 2025*. MIT NANDA. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

Chen, M., Tworek, J., Jun, H., Yuan, Q., Ponde de Oliveira Pinto, H., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., ... Zaremba, W. (2021). Evaluating large language models trained on code. *arXiv*. https://doi.org/10.48550/arXiv.2107.03374

Christensen, W., Sutton, J., & McIlwain, D. J. F. (2016). Cognition and the control of skilled action: Meshed control and the varieties of skill experience. *Mind & Language, 31*(1), 37–66. https://doi.org/10.1111/mila.12094

Dehghani, F., Dehghani, R., Naderzadeh Ardebili, Y., & Rahnamayan, S. (2025). Large language models in legal systems: A survey. *Humanities and Social Sciences Communications, 12*, Article

1977. https://doi.org/10.1057/s41599-025-05924-3

Dell'Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayer, 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

Denning, P. J., & Rousse, B. S. (2024). Can machines be in language? *Communications of the ACM, 67*(3), 32–35. https://doi.org/10.1145/3637629

Dreyfus, H. L. (2014). *Skillful coping: Essays on the phenomenology of everyday perception and action* (M. A. Wrathall, Ed.). Oxford University Press.

Dreyfus, H. L., & Dreyfus, S. E. (1988). *Mind over machine: The power of human intuition and expertise in the era of the computer*. Free Press. (Paperback Ed.)

Dreyfus, S. E., & Rousse, B. S. (2018). Commentary on Fernand Gobet's "The future of expertise: The need for a multidisciplinary approach." *Journal of Expertise, 1*(3), 181–183.

Ericsson, K. A. (2006). The influence of experience and deliberate practice on the development of superior expert performance. In K. A. Ericsson, N. Charness, R. R. Hoffman, & P. J. Feltovich (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 683–704). Cambridge University Press.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review, 100*(3), 363–406.

Everett, S. S., Lee, C. Y., & Xu, T. (2025). From tool to teammate: A randomized controlled trial of clinician-AI collaborative workflows for diagnosis. medRxiv. https://doi.org/10.1101/2025.06.07.25329176

Flores, F. (2012). *Conversations for action and collected essays: Instilling a culture of commitment in working relationships* (M. F. Letelier, Ed.). CreateSpace Independent Publishing Platform.

Gobet, F. (2015). *Understanding expertise: A multidisciplinary approach*. Palgrave Macmillan.

Gobet, F. (2018a). The future of expertise: The need for a multidisciplinary approach. *Journal of Expertise, 1(2),* 107–113.

Gobet, F. (2018b). Reply to Dreyfus and Rousse (2018). *Journal of Expertise*, 1(3), 184–186.

Gobet, F., & Chassy, P. (2009). Expertise and intuition: A tale of three theories. *Minds and Machines, 19*(2), 151–180. https://doi.org/10.1007/s11023-008-9131-5

Haugeland, J. (1998). Truth and rule-following. In *Having thought: Essays in the metaphysics of mind* (pp. 305–361). Harvard University Press.

Hosseini Maasoum, S. M., & Lichtinger, G. (2025). *Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data* (SSRN Scholarly Paper No. 5425555). Social Science Research Network. https://doi.org/10.2139/ssrn.5425555

Huising, R. (2019). Moving off the map: How knowledge of organizational operations empowers and alienates. *Organization Science, 30*(5), 1054–1075. https://doi.org/10.1287/orsc.2018.1277

Ingka Group. (2023, June 29). *AI and remote selling bring IKEA design expertise to the many.* https://www.ingka.com/newsroom/ai-and-remote-selling-bring-ikea-design-expertise-to-the-many/

Karpathy, A. (2026, April 30). Sequoia Ascent Summary. https://karpathy.bearblog.dev/sequoia-ascent-2026/

Klein, G. (1998). *Sources of power: How people make decisions*. MIT Press.

Lichtenberg, N. (2025, September 27). Accenture's $865 million reinvention includes saying goodbye to people without the right AI skills. *Fortune*. https://fortune.com/2025/09/27/accenture-865-million-reinvention-exiting-people-ai-skills/

MacIntyre, A. (2007). *After virtue: A study in moral theory* (3rd ed.). University of Notre Dame Press.

Mangiante, E. S., Peno, K., & Northup, J. (Eds.). (2021). *Teaching and learning for adult skill acquisition: Applying the Dreyfus and Dreyfus model in different fields*. Information Age Publishing.

McDowell, J. (2007). What myth? *Inquiry, 50*(4), 338–351. https://doi.org/10.1080/00201740701489211

Montero, B. G. (2016). *Thought in action: Expertise and the conscious mind*. Oxford University Press.

Niederhoffer, K., Kellerman, G. R., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (2025, September 22). AI-generated "workslop" is destroying productivity. *Harvard Business Review*. https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity

Reid, H. (2023, June 13). *IKEA bets on remote interior design as AI changes sales strategy.* Reuters. https://www.reuters.com/technology/ikea-bets-remote-interior-design-ai-changes-sales-strategy-2023-06-13/

Rousse, B. S. (2019). Self-awareness and self-understanding. *European Journal of Philosophy*, *27*(1), 162–186. https://doi.org/10.1111/ejop.12377

Rousse, B. S. (2026a). Toward criteria for artificial self-consciousness: Unity, normativity, and agency. *Proceedings of the AAAI Symposium Series*, 8(1), 335–344. https://doi.org/10.1609/aaaiss.v8i1.42563

Rousse, B. S. (2026b). Heidegger and phenomenological approaches to work. In J. D. Jonker & G. Rozeboom (Eds.), *The Oxford handbook of the philosophy of work*. Oxford University Press.

Rousse, B. S. (in press). Care, human enfeeblement, and the existential implications of AGI. In M. Iklé, A. Franz, & A. Kemendo (Eds.), Artificial general intelligence: 19th International Conference, AGI 2026, San Francisco, CA, USA, July 27–30, 2026, proceedings. Springer.

Rousse, B. S., & Dreyfus, S. E. (2021). Revisiting the six stages of skill acquisition. In E. Mangiante, K. Peno, & J. Northup (Eds.), *Teaching and learning for adult skill acquisition: Applying the Dreyfus and Dreyfus model in different fields* (pp. 3–28). Information Age Publishing.

Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.

Searle, J. R. (1979). A taxonomy of illocutionary acts. In *Expression and meaning: Studies in the theory of speech acts* (pp. 1–29). Cambridge University Press.

Simon, L. K. (2025, August 4). Is AI responsible for the rise in entry-level unemployment? AI drives a wedge between junior and senior hiring. Revelio Labs. https://www.reveliolabs.com/blog/ai-drives-a-wedge-between-junior-and-senior-hiring.

Smith, B. C. (2019). *The promise of artificial intelligence: Reckoning and judgment*. MIT Press.

Sudnow, D. (2001). *Ways of the hand: A rewritten account*. MIT Press.

Sutton, J. (2007). Batting, habit and memory: The embodied mind and the nature of skill. *Sport in Society, 10*(5), 763–786. https://doi.org/10.1080/17430430701442462

Taylor, C. (1995). To follow a rule. In *Philosophical arguments* (pp. 165–180). Harvard University Press.

Winograd, T., & Flores, F. (1986). *Understanding computers and cognition: A new foundation for design*. Ablex Publishing Corporation.

Wrathall, M. (2017). Introduction: Background practices and understandings of being. In H. L. Dreyfus, *Background practices: Essays on the understanding of being* (M. A. Wrathall, Ed., pp. 1–18). Oxford University Press.

Received: November 18, 2025
Revision received: May 20, 2026
Accepted: June 6, 2026