AI Skills for College Graduates

Exploring How Instructors and Employers Prioritize AI Skills Differently

Michael Fried

August 18, 2026

ITHAKA S+R

Through research, collaboration, and innovation, Ithaka S+R informs policy, guides strategies, and builds evidence-based tools to broaden access to postsecondary education, improve student and workforce outcomes, and advance scholarship.

Ithaka S+R is part of ITHAKA, a nonprofit with a mission to improve access to knowledge and education for people around the world. We believe education is key to the wellbeing of individuals and society, and we work to make it more effective and affordable.

We would like to thank Dr. Kenneth Matos, director of research insights at HiBob, for creating and sharing the AI Skills Framework as well as his research with employers.

ITHAKA S+R

Introduction

Since the public release of ChatGPT in 2022, both higher education institutions and employers have been keenly focused on how to prepare college students, and ultimately graduates, for the new world of AI-inflected work.[1] Many colleges and universities are jockeying to be among the first to provide access to enterprise AI platforms or integrate AI into their curricula, despite some student and instructor misgivings.[2] Local, state, and federal policymakers are also becoming involved in providing guidance on how educational institutions should prepare students for work in the age of AI.[3]

But what does it mean to be ready for the AI workplace? While there is enormous collective attention on students' and workers' skills with using AI, it is less clear what "AI skills" should encompass. [4] AI skills are as new to employers as they are to graduating students, and the technology itself is changing so quickly that many may struggle to remain on the cutting edge. Establishing a specific, actionable, and assessable framework for what constitutes AI skills is a necessary first step in aligning efforts to sufficiently support students as they transition from higher education to employment in the age of AI.

To that end, we embarked on one of the first major efforts to vet a comprehensive framework of what constitutes AI skills with college and university instructors as well as to understand how instructors prioritize these skills. The framework comes from HiBob, a multi-national human

---

[1] "How Universities Can Prepare Students for AI-First Workplaces," *Boston Consulting Group*, updated February 4, 2026, https://www.bcg.com/publications/2026/preparing-students-ai-first-workplaces.

[2] Joshua Bay, "How 5 Colleges Are Approaching AI," *Inside Higher Ed*, April 3, 2026, https://www.insidehighered.com/news/student-success/academic-life/2026/04/03/how-5-colleges-are-approaching-ai.

[3] Erin Whinnery, Adrienne Fischer, and Heena Kuwayama, *State Artificial Intelligence (AI) Guidelines* (Denver, CO: Education Commission of the States, April 2026), response to information request.

[4] Scott Carlson, "Can Colleges Make All Their Students 'AI Fluent'?" *The Chronicle of Higher Education*, June 17, 2026, https://www.chronicle.com/article/can-colleges-make-all-their-students-ai-fluent; Elisa J. Sobo and David M. Goldberg, "What Do Employers Mean by 'AI Skills,' Anyway?" *Inside Higher Ed*, April 10, 2026, https://www.insidehighered.com/opinion/views/2026/04/10/what-do-employers-mean-ai-skills-anyway-opinion.

resources technology company that has used this framework with employers.[5] By combining data on how employers and instructors view the skills in the framework we can identify points of alignment and dissonance. If institutions or instructors want to better align their curricula, assessments, or learning outcomes to employer priorities, this analysis may prove useful.

AI Skills Framework

The AI Skills Framework was originally developed by researchers at HiBob as part of an effort to better understand what specific AI-related skills employers seek and might be willing to pay a premium for when hiring. The framework consists of seven broad categories:

- AI literacy
- Continuous learning orientation
- Prompting and input quality
- Evaluating and improving output quality
- AI safety, ethics, and governance
- Workflow evaluation and redesign
- Automation and technical integration

These categories are meant to incorporate the majority of how a typical, non-technical, employee would use AI tools in the workplace, whether those tools are foundational LLM chatbots or more purpose-built systems that leverage LLM models. The framework also includes a set of 26 associated skills; the full list of skills along with their definitions are included in Appendix A.

[5] "The AI Skills Report 2026," *HiBob*, 2026, https://www.hibob.com/lp/ai-transformation/.

Methodology

To reach a target sample of 500 participants, we invited 28,200 instructors from four-year, nonprofit colleges and universities to respond to the survey. We sent the first email to this group on April 13, 2026, and sent one reminder to 22,000 of those invitees approximately one week following the initial invite. Data collection concluded on April 24, 2026.

We asked respondents to rate the importance of 26 specific AI-related skills for graduates of their institutions and whether they personally taught those skills in their courses. Respondents rated each skill on a seven-point scale from "Not at All Important" to "Essential/Critical." The survey also included open-response opportunities as well as questions about respondents' personal use of AI in their teaching and how their institution approaches AI policy and practice. All respondents were actively teaching undergraduate students during the Spring 2026 academic term.

Separately, researchers at HiBob surveyed a multi-national panel of employers on how important proficiency in each of the 26 skills is for AI users withing their organizations. The HiBob survey was completed by 1,200 employers in February and March 2026, 200 of which operate in the United States.[6] It is this smaller, US-based sample against which the instructor responses to Ithaka S+R's survey are compared.

Because we closed the survey when we reached a pre-determined total of 500 respondents, it may be that the data includes early response bias. It is possible for survey respondents to differ based on when they take a survey; for example, survey respondents who respond early to a survey versus those who respond later (perhaps after several weeks or with multiple reminders) may be more interested or engaged with the survey topic, which may impact the findings.

To minimize the time and effort for participants in completing the survey, this report relies on individual data included in the email list we purchased from a marketing agency for some data points, like institutional characteristics and discipline. This approach results in some limitations on

---

[6] More information about the HiBob survey, methodology, and results can be found at https://www.hibob.com/research/ai-skills-2026-report

the reporting on institutional information and disciplinary affiliation, which is limited by the data structures of the sample provider.

The majority of respondents self-reported as tenured or tenure-track instructors, with a substantial plurality indicating they held an adjunct or contingent position. The instructors who indicated some other kind of role mostly serve in long-term contract instructional roles that do not include tenure status.

Table 1

Instructor Status	Number of Respondents
Tenured	214
Adjunct/contingent	122
Tenure-track	87
Other	71
Graduate student instructor	5

Almost three-fifths of the participants are employed at doctoral institutions, and fewer than 12 percent teach at exclusively undergraduate colleges. All participants indicated that they were presently teaching at least one undergraduate course at the time of the survey. More than three-fifths of participants work at public institutions while the remainder work at private, not-for-profit institutions.

Table 2

Institutional Type	Number of Respondents
Doctoral	299
Masters	132
Baccalaureate	58
Special Focus	11
Public	315
Private, Not-for-Profit	185

The participants represent a wide range of disciplines and academic departments. The plurality comes from the social and natural sciences as well as business, with other disciplinary groupings containing 20 or fewer respondents.

Table 3

Academic Discipline	Number of Respondents
Social Sciences	76
Natural Sciences	53
Business	51
Education	39
Fine Arts	36
English Literature	33
Interdisciplinary Studies	28
Computer Science	27
Engineering	25

Communications	23

Additional 11 rows not shown.

Notable resistance to AI

Several recipients of the survey invitation email responded directly to express a negative opinion about AI in general, AI in higher education, or the survey itself. These responses expressed various levels of pessimism about or disapproval of AI along with the message that they would not be participating in this research effort because of those beliefs. It is safe to assume that some unknown number of potential respondents declined to participate on similar grounds but chose not to reply to the invitation email. Because non-participation in this research is the defining behavior for some subset of instructors who have strong objections to AI, it can be assumed that the results reported throughout include some degree of a positive bias.

Some participants echoed these negative sentiments in their narrative responses to the survey, meaning that these perspectives are captured to some degree in the data and analysis presented here. For example, one instructor suggested that knowing "[h]ow to disable and resist AI is the most important skill." Another offered that the "skill of doing things without AI" should be prioritized for college graduates. The survey did not directly ask respondents about their opinion of AI in higher education, but it is clear from the responses to survey invitation and the open-text questions that there is strong resistance from many instructors.

Skills alignment

In this section we compare the degree to which instructors value each of the skills within the framework to the priority US-based employers ascribe to those skills. Across the 26 skills, instructors and employers agreed, on average, on the importance of only one skill: setting realistic expectations for AI-augmented work. In most other cases, one group or the other believed the various skills to be of greater or lesser import. How these differences changed between the skill categories reveal the diverging priorities for the two groups.

Instructors place the highest priority on the responsible use of AI (uses AI responsibly) along with several skills relating to navigating the limits of both human and AI capabilities and contributions (reinforces human accountability, transparency and attribution, proactively reviews output quality, and recognizes the limits of their own expertise). One instructor noted, "…understanding the limitations of AI and the fact that LLMs cannot replace human intelligence is key." Collectively, the skills that instructors prioritize a critical approach to using AI wherein understanding the limits of its capabilities and outputs are paramount. These priorities align with many foundational academic values, including appropriate attribution, review and revision of work, and information literacy.

> Instructors place the highest priority on the responsible use of AI (uses AI responsibly) along with several skills relating to navigating the limits of both human and AI capabilities and contributions.

Conversely, employers are more concerned with skills related to workflows, automation, and how people and AI can work together efficiently. Employers prioritize skills more related to the activities of workers than students, which may account for the difference in prioritization between employers and instructors. Further, employers highly value human-to-human and human-to-machine skills more than instructors, which may be a reflection on the more team-oriented nature of workplaces where efficiency is prized in contrast to the often-individualistic nature of evaluation in academic settings where domain knowledge and critical thinking prevail.

> Employers highly value human-to-human and human-to-machine skills more than instructors.

The chart below compares the average prioritization of each skill by instructors and employers on a seven-point scale:

Chart 1: AI Skill Priorities

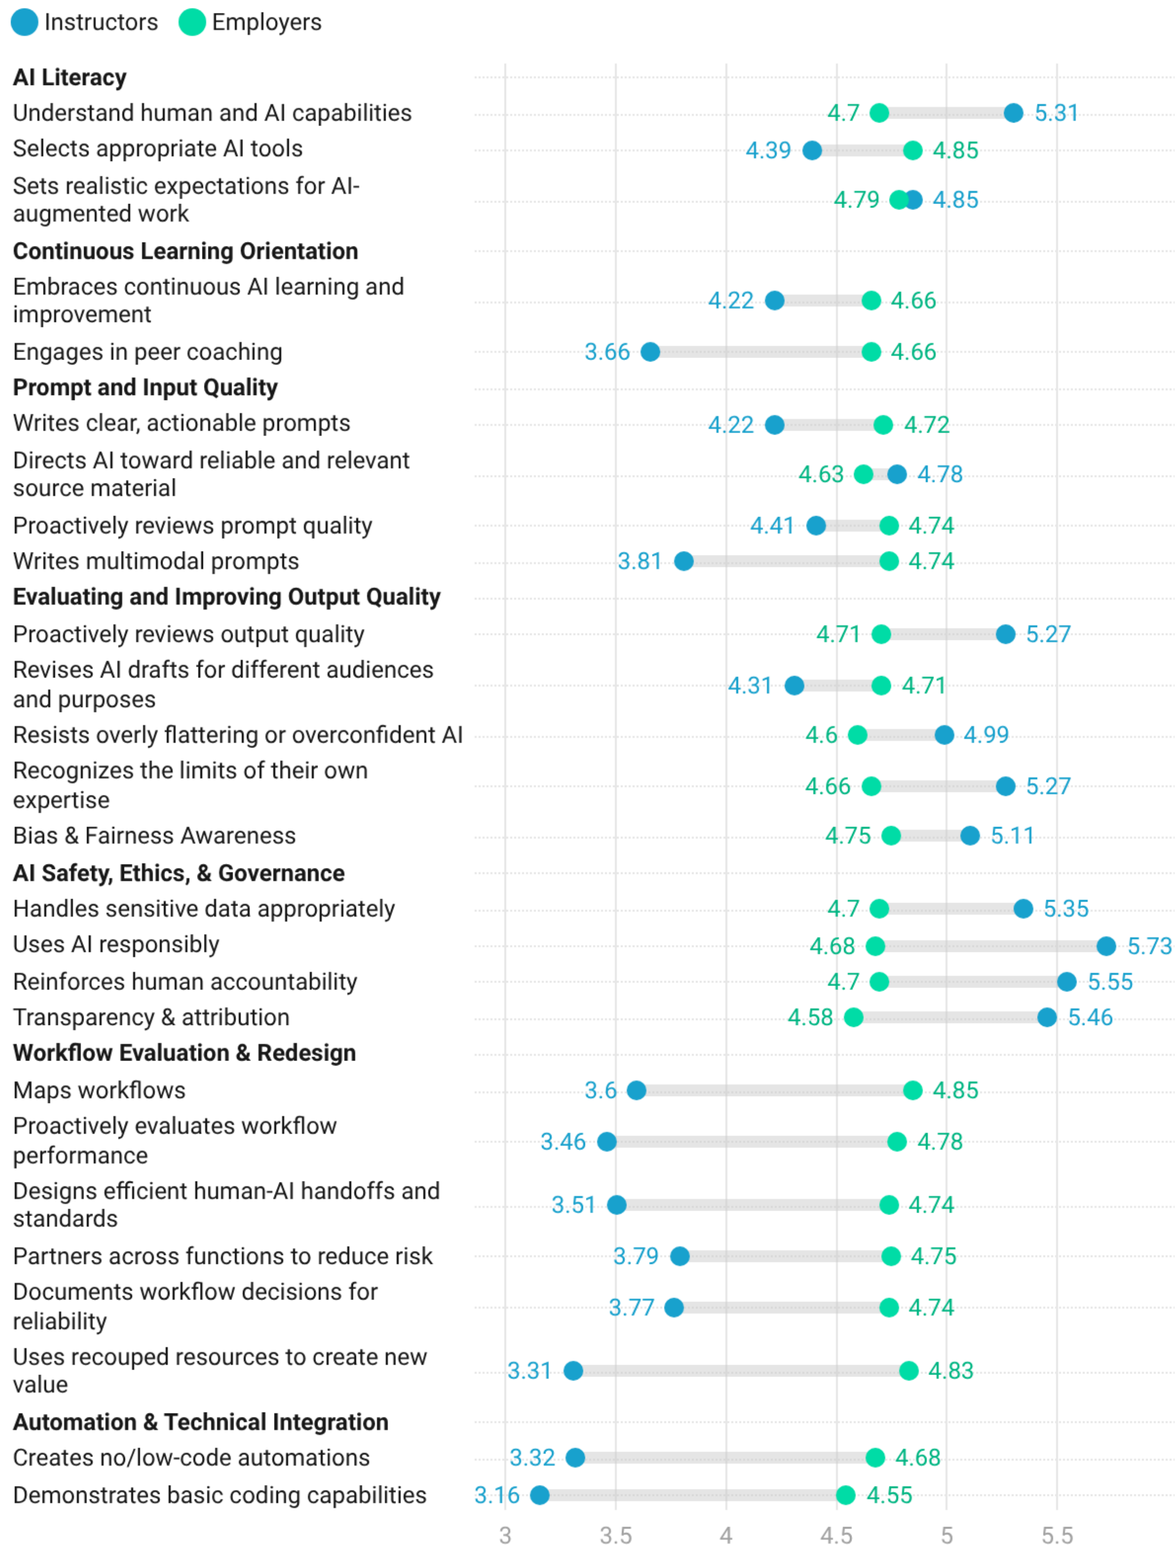

Created with Datawrapper

Teaching practices

The survey also included questions about whether instructors personally taught each of the skills in their courses. Two-thirds of instructors teach the transparency and attribution skill, likely in the context of academic integrity and scholarly attribution.[7] Roughly half of the respondents also teach skills related to the responsible and cautious use of AI and its outputs (responsible use of AI, proactively reviews output quality, reinforces human accountability), reflecting the value that academia places on critical thinking and centering the human role in AI use. One instructor suggested teaching an "approach to AI that distinguishes the dignity and inherent worth of humans over their tools."

> There are whole categories of skills which employers value that instructors do not, and likely as a result, do not teach in their courses.

Far fewer instructors taught skills in the workflow evaluation and redesign and technical integration categories. These are also the areas where instructor and employer prioritization had the largest divergence. This dissonance is the strongest evidence that there is an AI skills gap between higher education and employers. There are whole categories of skills which employers value that instructors do not, and likely as a result, do not teach in their courses. One instructor spoke directly to this difference in priorities, saying, "...many of these [survey] questions stress workflow and productivity, I would stress to my students the importance of fact-checking and cross-referencing."

The table below shows the percentage of instructors who report teaching each of the AI-related skills, organized by skill category.

---

[7] Clay Shirky, "The Post-Plagiarism University," *The Chronicle of Higher Education*, November 3, 2025, https://www.chronicle.com/article/the-post-plagiarism-university.

Table 4: AI Skills Percent Teaching

AI Literacy	
Understand human and AI capabilities	50.0%
Selects appropriate AI tools	28.6%
Sets realistic expectations for AI-augmented work	40.6%
Continuous Learning Orientation	
Embraces continuous AI learning and improvement	24.8%
Engages in peer coaching	18.4%
Prompt and Input Quality	
Writes clear, actionable prompts	36.2%
Directs AI toward reliable and relevant source material	33.6%
Proactively reviews prompt quality	23.2%
Writes multimodal prompts	15.8%
Evaluating and Improving Output Quality	
Proactively reviews output quality	49.2%
Revises AI drafts for different audiences and purposes	30.2%
Resists overly flattering or overconfident AI	34.0%
Recognizes the limits of their own expertise	43.8%
Bias & fairness awareness	43.8%
AI Safety, Ethics, & Governance	
Handles sensitive data appropriately	35.8%
Uses AI responsibly	51.4%
Reinforces human accountability	49.2%
Transparency & attribution	66.2%
Workflow Evaluation & Redesign	
Maps workflows	16.0%
Proactively evaluates workflow performance	11.6%
Designs efficient human-AI handoffs and standards	10.6%
Partners across functions to reduce risk	11.8%
Documents workflow decisions for reliability	14.0%
Uses recouped resources to create new value	8.2%
Automation & Technical Integration	
Creates no/low-code automations	10.2%
Demonstrates basic coding capabilities	11.0%

Created with Datawrapper

Institutional context for AI skills

The Ithaka S+R survey included questions about instructors' view of the institutional context in which they are teaching. The responses echo other research that suggest most US colleges and universities have not developed coherent AI strategies and policies.[8] Although respondents slightly disagreed that their institutions expect undergraduate students to acquire a moderate level of AI proficiency before completing a degree, they more strongly disagreed that their institutions had a consensus about what AI skills actually looked like. One instructor noted that their "institution announced AI literacy as a top priority—and then did nothing since," leaving them feeling "100% unqualified teaching anything about AI" to their students.

Similarly, respondents strongly disagreed that institutions had a shared framework for assessing students' AI skills. A shared framework to assess AI skills is a prerequisite for institutions to be able to determine the degree to which their students have acquired a moderate level of AI proficiency. These findings together indicate that instructors do not believe institutions are fully prepared to respond to AI. As one instructor put it, "I haven't changed my teaching much because my institution is still very anti-AI, but I listed many of these skills as critically important (even though I don't teach about them)." Instructors responded on a five-point agreement scale, from 1- Strongly disagree to 5 – Strongly agree.

[8] For example, see Erin McCusker and Russell Michalak, "AI Policies in US Universities: A Critical Analysis of Policy Gaps and Library Involvement," *Journal of Library Administration* 65, nos. 6–7 (2025): 808–824; Claire Baytas and Dylan Ruediger, "Making AI Generative for Higher Education: Adoption and Challenges Among Instructors and Researchers," *Ithaka S+R*, May 1, 2025,https://doi.org/10.18665/sr.322677.

Chart 2: Expectations Exceed Infrastructure on AI Skills

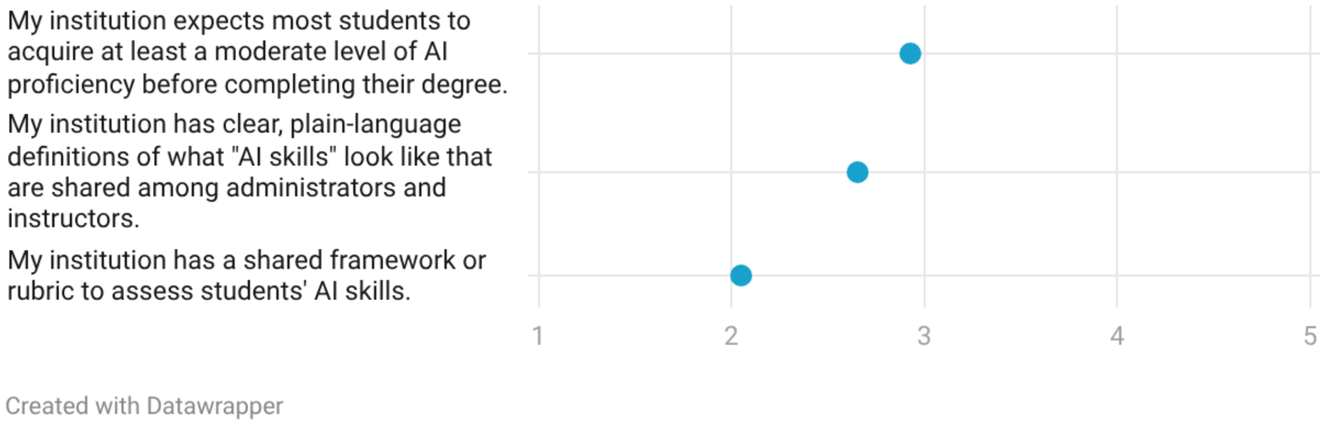

Discussion

The results of these two survey initiatives illustrate areas where instructors and employers differ and align in their priorities around AI-related skills for recent college graduates. Instructors and employers are in closest agreement around the AI literacy skill of setting realistic expectations for AI-augmented work and directing AI toward reliable and relevant source material. There is the highest level of disagreement in the workflow evaluation and redesign and automation and technical integration skill categories, with instructors prioritizing these areas well below employers.

These results also show that institutions and instructors are in the early stages of organizing around AI-ready learning environments. Few instructors report that their institutions have built consensus around what AI-related skill outcomes are important for their students and even fewer report that their institutions have provided the necessary assessment apparatus to know if students are reaching those outcomes. Of the AI-skills included in this survey, only three skills are taught by half or more instructors, possibly leaving many skills out of degree programs. Taken together, these findings point to opportunities for institutions seeking greater alignment with employers, as well as for developing more sophisticated approaches to preparing students for the ubiquity of AI during and after college.

The data and analyses in this report are broad by design to offer an initial snapshot alignment between instructors and employers, leaving more nuanced lines of inquiry for further research. Exploring how these findings may differ between disciplines will be an important step in understanding how the broad disciplinary groupings (liberal arts, sciences, social sciences, arts, humanities, etc.) and individual degree programs respond to AI. This report also did not include instructors from community colleges or vocational institutions, which often have a different kind of relationship with employers than four-year institutions. Further research should address how these parts of higher education prioritize and implement AI-related skill instruction. The larger question of to what degree, if any, there should be alignment between the skills instructors teach and what employers look for in new hires is beyond the scope of this report, though this "pulse check" may contribute to that evolving conversation.

Conclusion

As a phrase, "AI skills" is an amorphous category into which many different attitudes, capacities, and understandings can fit. As conversations about generative AI continue between and about instructors, employers, and students, it is imperative that the vagaries around AI skills distill into a more discrete and concrete set of competencies that can be integrated into college curricula and job descriptions. This research aims to move both corporate organizations and higher education institutions toward that greater specificity by exploring how employers and instructors differently prioritize a specific set of AI-related skills.

> There is meaningful, though likely not insurmountable, dissonance about the importance instructors and employers attribute to specific AI-related skills.

There is meaningful, though likely not insurmountable, dissonance about the importance instructors and employers attribute to specific AI-related skills. The framework presented here can serve as part of a first step in bridging those gaps by offering an actionable set of specific AI-related skills for employers, institutions, and instructors to use when making strategic decisions about how AI will manifest in their respective jobs, curricula, and classrooms. Our intent is not to dictate what skills should—or should not—be prioritized. Rather, it is to bring the conversation about AI-skills out of the realm of vaporware and into real conversations and reflections that instructors, administrators, and employers are having every day.

Appendix A: The AI Skills Framework

AI Literacy
Understands human and AI capabilities — They know which tasks are best performed by AI or a person based on task characteristics (e.g. risk, reliability, required judgment, speed, etc.) and AI and human strengths (e.g. speed vs. judgement).
Selects appropriate AI tools — They pick the right AI tool or function (e.g. basic or advanced models) for the job (e.g. drafting, summarizing, analysis, search, automation) and can explain why it fits.
Sets realistic expectations for AI-augmented work — They can clearly explain AI enhancement (e.g., drafting) versus what still takes human attention and time (e.g., reviews, approvals), so timelines and outcomes stay realistic.

Continuous Learning Orientation
Embraces continuous AI learning and improvement — They regularly reassess evolving AI tools to update prompts, agents, and automations to utilize new capabilities. They take an iterative, trial-and-error approach to AI work — testing changes, learning from results, and refining until the outcome consistently meets the intended standard.
Engages in peer coaching — They both give and seek coaching to standardize AI best practices, observing how others work, exchanging specific feedback on outputs, sharing prompts and examples and building repeatable routines that improve individual and collective performance over time.

Prompting and Input Quality
Writes clear, actionable prompts — They write prompts providing guidance for how the AI should operate (e.g. role/context, constraints, examples, etc.) and expectations for the output (e.g., length, tone, reading level, etc.) to improve the final output.
Directs AI towards reliable and relevant source material — They ground AI outputs in credible sources and highlighting the most important information from messy notes or documents with priorities, constraints, and exclusions so the AI produces a reliable, relevant response.
Proactively reviews prompt quality — They test prompts on complex or sensitive cases and add rules to prevent common failures like made-up details, privacy slips, or misleadingly confident answers.

Writes multimodal prompts — They can effectively craft prompts using multiple forms of data and media (e.g. documents, images, tables, video, audio, etc.) and specify how outputs should be extracted, summarized, or validated.

Evaluating and Improving Output Quality
Proactively reviews output quality — They verify key claims across multiple sources, ensuring correct figures (e.g. totals, ratios, or trends) and appropriate clarity, tone, and depth in text.
Revises AI drafts for different audiences and purposes — They revise AI drafts into audience and goal-specific products, adjusting message, detail level, tone, and accessibility elements (e.g., headings, plain language, alt text, etc.) while keeping facts accurate and consistent.
Resists overly flattering or overconfident AI — They don't accept confident or flattering AI outputs without evidence, they surface uncertainty, and they prevent tone from driving decisions.
Recognizes the limits of their own expertise — They can identify high-stakes or unfamiliar situations where their knowledge and skills are insufficient to properly evaluate the output and bring in subject-matter experts to verify their evaluations.
Bias & fairness awareness — They recognize biased framing or unequal treatment in AI outputs and take corrective steps (e.g., re-prompt, add constraints, escalate, etc.).

AI Safety, Ethics, & Governance
Handles sensitive data appropriately — They recognize personal/confidential data, anonymize or avoid it, and follow approved tools/channels and retention rules.
Uses AI responsibly — They follow ethical, legal, and company standards (e.g., protecting private and owned information) when using AI, especially for regulated or high-risk topics such as finance, health, or legal matters.
Reinforces human accountability — They make sure a person is clearly responsible for decisions made with AI support, define which steps require human review or sign-off, and ensure AI outputs are treated as recommendations—not final decisions.
Transparency & attribution — They appropriately disclose when AI meaningfully helped content creation or analysis and distinguish AI-generated content from human conclusions.

Workflow Evaluation & Redesign
Maps workflows — They capture how existing workflows actually operate end-to-end (steps, roles, tools, inputs/outputs, decision points, exceptions, and handoffs) in a clear format that others can follow, validate, and use as a shared baseline.
Proactively evaluates workflow performance — They assess workflows using practical indicators (e.g., effort, cost, collaboration quality, etc.), identifying delay and failure points, and opportunities to recoup resources (e.g., time, effort, materials, etc.).

Designs efficient human–AI handoffs and standards — They help people and AI work together smoothly by defining new inputs, outputs, roles, approval points, quality standards, and where work is stored or routed—to minimize stalls, duplication, and confusion.
Partners across functions to reduce risk — They involve the right partners and subject matter experts (e.g., Legal, Security, IT, etc.) early, bringing clear context and examples, and incorporating cross-functional requirements into the workflow design and documentation.
Documents workflow decisions for reliability — They keep clear, up-to-date process notes on AI inputs, steps, and outputs so others can follow the workflow, issues can be resolved quickly, and reviews for compliance or audits are straightforward.
Uses recouped resources to create new value — They adjust activities to direct savings towards higher value activity (e.g., reducing backlog, product innovation, stakeholder relations, etc.) and can point to the tangible operational effect of their changes (e.g., faster service, higher-quality deliverables and stakeholder relationships).

Automation & Technical Integration
Creates no/low-code automations — They create simple, reliable automations that connect AI to everyday tools (e.g., email, spreadsheets, chats, forms, etc.) to reduce manual and repetitive steps.
Demonstrates basic coding capabilities — They can use basic APIs (application programming interfaces) or light scripting to connect systems and troubleshoot problems. They can read simple API documentation, use keys securely, handle common failures (timeouts, missing fields, permission errors), respect rate limits, and add simple logging so workflows can be monitored and fixed without guesswork.