ISSN: 2473-4901
v12 n6753

AI Literacy in the Wild: What Anthropic, OpenAI, and Google Teach in Their AI Academies

Guido Lang
guido.lang@qu.edu

Luis SaCouto
luis.sacouto@qu.edu

Vivian Etim
vivian.etim@qu.edu

Anthony Asare
anthony.asare@qu.edu

Quinnipiac University
Hamden, CT, USA

Jason Sharp
jason.sharp@hsutx.edu
Hardin-Simmons University
Abilene, TX 79698

Abstract

AI literacy is an important concept that has been defined as the ability to understand, use, evaluate, and act responsibly with AI. Leading AI companies such as Anthropic, OpenAI, and Google have built large public AI academies that train users on how to use current AI tools. This paper studies those AI academies and compares them with academic work on the AI literacy, in order to develop an integrated framework that builds on both practitioner and academic frameworks. Using an inductive analysis of lessons from Anthropic Academy, OpenAI Academy, and Google AI Academy conducted in May 2026, the authors derive a three-tier, nine-dimension framework of industry AI literacy. The authors then map that framework to recent academic work on AI literacy, using the AI Literacy Heptagon as the main organizing reference. The mapping shows strong agreement around critical evaluation, responsible use, and workflow integration. It also shows that the industry academies teach several more concrete skills that are treated only at a general level in the academic literature, including context management, reusable AI artifacts, agent supervision, and evaluation engineering. The paper argues that academic and industry views of AI literacy are not opposed, because academic work asks what an AI-literate person should understand while industry training asks what an AI-literate person can do and remain responsible for.

Keywords: AI literacy, generative AI, AI education, AI academies, higher education, competency framework

ISSN: 2473-4901
v12 n6753

AI Literacy in the Wild: What Anthropic, OpenAI, and Google Teach in Their AI Academies

Guido Lang, Luis SaCouto, Vivian Etim, Anthony Asare, and Jason Sharp

1. INTRODUCTION

The meaning of AI literacy has changed quickly since the release of widely used generative AI tools. Early AI literacy work often asked whether users could recognize AI, understand basic machine learning ideas, and judge the social and ethical effects of AI systems. Those questions remain important, although everyday AI use now also asks people to write prompts, manage context, choose tools, check outputs, build reusable workflows, and supervise AI agents that can act across files, websites, and software tools.

Academic research has responded to this change by moving beyond basic awareness toward practical use, ethical reasoning, social impact, legal knowledge, and the integration of AI into work and learning. For example, Hackl, Mueller, and Sailer (2026) review 27 AI literacy sources and group the higher-education literature into seven dimensions, showing that the field has become broader and more applied while still treating AI literacy as an educational and civic construct.

At the same time, the major AI providers have created their own AI literacy programs. Anthropic Academy, OpenAI Academy, and Google AI Academy are important educational resources because they show what the firms that build leading AI systems believe users need to learn. These academies include basic AI use, prompting, responsible use, workplace tasks, app building, agents, tool use, and evaluation. Their content gives a practical view of AI literacy that can be studied inductively before imposing categories from the academic literature.

The academies also provide evidence of current industry priorities because each provider decides which capabilities, workflows, responsibilities, and technical practices deserve formal training. For information systems education, these choices offer a timely signal about the competencies being expected in AI-enabled organizations. For IS research, they provide an observable record of how major technology firms translate rapidly changing AI capabilities into roles and practices. Provider curricula should still be interpreted critically because they also reflect product strategies and may give less attention to social, legal, and regulatory concerns.

This study aims to address two research questions:
1. What framework of AI literacy emerges from the public academy content of Anthropic, OpenAI, and Google?
2. How does that industry-derived framework map to recent academic AI literacy work?

A three-tier, nine-dimension framework is proposed in which the Fundamentals tier covers the skills expected of an AI user, the Professional tier covers workplace use, workflow design, and agent supervision, and the Builder tier covers the design and testing of AI systems. Rather than replacing academic AI literacy work, the framework adds more detail about present-day work with AI tools, especially in the middle area between casual user and software developer.

2. RELATED LITERATURE

Hackl, Mueller, and Sailer (2026) provide the most useful academic background for this study because the authors define AI literacy in a way that supports curriculum design in higher education. The authors begin from the problem that AI literacy has become a broad and uneven field. Earlier work often stressed technical understanding, such as knowing how AI systems process data, make predictions, or use algorithms. More recent work adds practical use, ethical reasoning, social effects, and professional application. Hackl and colleagues argue that higher education needs a framework that can hold these

ISSN: 2473-4901
v12 n6753

parts together without reducing AI literacy to either technical skill or general awareness. Their paper therefore reviews the current AI literacy literature, identifies gaps in existing definitions, proposes a structured model, and tests whether that model can be used to map academic curricula.

Long and Magerko's (2020) framework first defined AI literacy through competencies such as recognizing AI, understanding what AI can and cannot do, understanding basic representation and decision-making, and thinking about ethics and human roles. Ng et al. (2021) extend this work through a layered model that includes knowing and understanding AI, using and applying AI, evaluating and creating with AI, and AI ethics. These older frameworks help establish AI literacy as more than tool use, although they were developed before the current wave of generative AI tools became common. Hackl et al. (2026) also include more recent work on higher education, generative AI, and teacher or student competencies. For example, the EDUCAUSE framework (Kassorla, Georgieva, & Papini, 2024) is presented as a practical higher-education approach that stresses teaching, learning, critical review of AI tools, and attention to bias and misuse. UNESCO's AI competency framework (Miao, F., & Cukurova, M. (2024)) is treated as a broader policy framework that emphasizes understanding, evaluation, responsible use, creation, inclusion, and equity. The OECD and European Commission AILit framework (2025), although aimed at primary and secondary education, is important because it organizes AI literacy around the four domains of engaging with AI, creating with AI, managing AI, and designing AI.

Across the reviewed literature, Hackl et al. (2026) identify both progress and fragmentation, arguing that current AI literacy definitions have moved from narrow technical knowledge toward a more complete view that includes social, ethical, legal, and applied dimensions. At the same time, they find that the field still lacks a unified higher-education definition and an actionable framework for implementation. Many frameworks name important competencies but do not give educators a simple way to compare programs, identify gaps, or choose target levels for different learners. This is especially important in higher education because AI literacy is not the same for all students. A media pedagogy student, business student, engineering student, and computer science student may all need AI literacy, although they may not need the same depth in every dimension. Hackl et al. therefore, frame AI literacy as both generic and domain-specific, with all students needing a baseline while some programs require higher levels in technical, professional, or design-oriented areas.

Hackl et al. (2026) review of 27 AI literacy frameworks shows which AI literacy components appear most often in the literature. Application proficiency appears in 26 of 27 frameworks, technical knowledge and skills in 24, critical thinking ability in 23, and ethical awareness and reasoning in 20. Integration skills appear in 18 frameworks, social impact understanding in 11, and legal and regulatory knowledge in only 2. These counts are important because the authors do not simply keep the most common categories. Instead, they use the review to combine recurring dimensions with emerging dimensions that higher education will need, including legal and regulatory knowledge even though it appears in only a small number of sources. The authors include this dimension because AI-related law, copyright, data protection, and regulation are becoming part of responsible AI use. They also treat integration skills as central because AI is increasingly embedded in professional and academic workflows rather than used only as a stand-alone tool.

The resulting AI Literacy Heptagon shown in Figure 1 contains seven dimensions that together describe technical understanding, applied use, critical judgment, ethical reasoning, social effects, workflow integration, and legal or regulatory knowledge. **Technical Knowledge and Skills** refers to understanding basic AI principles, data processing, algorithms, computational mechanisms, model limits, and the ways AI systems generate outputs. **Application Proficiency** refers to using AI tools across contexts, including choosing tools, developing prompts, improving work processes, assessing outputs, and adapting AI applications to a task or domain. **Critical Thinking Ability** refers to judging AI systems, claims, outputs, limits, and use cases rather than accepting results without evidence. **Ethical Awareness and Reasoning** covers bias, privacy, safety, transparency, accountability, autonomy, fairness, labor effects, democratic concerns, human relationships, and environmental costs. **Social Impact Understanding** focuses on the broader effects of AI on work, education, institutions, social structures, and different populations. **Integration Skills** refers to incorporating AI into digital environments and workflows, identifying where AI fits, adapting work processes, and managing shifts between human and AI-driven work. **Legal and Regulatory Knowledge** covers AI-specific rules, data

 ISSN: 2473-4901
v12 n6753

protection, copyright, regional differences in regulation, and the fact that existing legal regimes still apply to generative AI (Hackl et al., 2026).

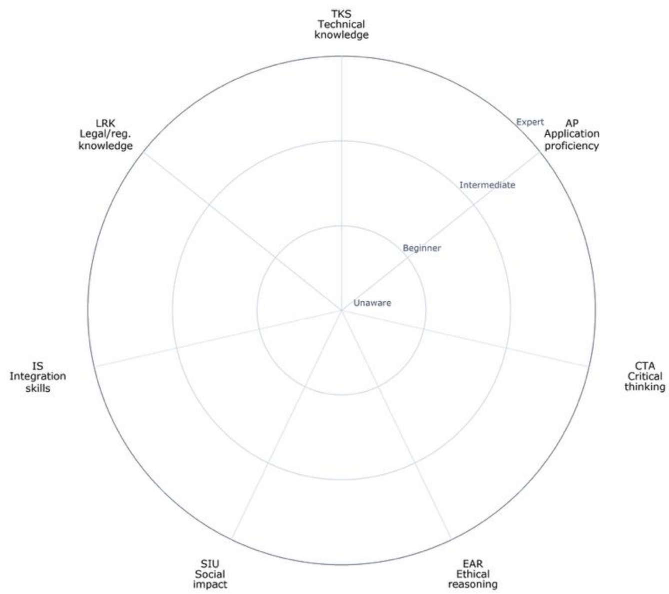

Figure 1. AI Literacy Heptagon (Hackl et al., 2026)

A key contribution of Hackl et al. (2026) is that the Heptagon is designed for curriculum mapping as well as conceptual discussion. The authors add the four proficiency levels of Unaware, Beginner, Intermediate, and Expert, which lets educators describe both breadth and depth. A general education curriculum might aim for beginner-level competence across all seven dimensions, while an AI Engineering curriculum might aim for Intermediate or Expert levels in technical knowledge, application proficiency, integration, and evaluation-related work. The authors test this idea through expert-led curriculum mapping in the AI Engineering and Media Pedagogy programs. These examples show that the same framework can produce different profiles depending on program goals. In the context of this study, the Heptagon provides a structured way to map an inductively derived industry framework onto the academic curriculum-design language of higher education.

3. METHODOLOGY

Publicly-available AI literacy courses from Anthropic Academy, OpenAI Academy, and Google AI Academy were collected and analyzed in May 2026, with each lesson serving as the unit of analysis. Courses were retained when they taught general or professional AI competencies and excluded when designed primarily for a specific sector or audience. The Anthropic corpus covers AI capabilities, prompting, computer use, agents, and related topics, while excluding sector-specific offerings such as AI for educators or nonprofits. The OpenAI corpus covers the Work Users and Builders tracks, and the Google corpus covers the Google AI Essentials Specialization and Google AI Professional Certificate.

Across Anthropic and OpenAI, the corpus contains 469 lessons, including 210 with fully extractable text and 259 video lessons without transcripts for which only lesson titles and surrounding course structure were available. The 193 Google lessons all had fully extractable text, bringing the combined corpus to 662 lessons. The analysis preserved this difference in source granularity by using text-rich lessons for detailed competency summaries and limiting conclusions from title-only lessons to conservative inferences grounded in their titles and course positions, using source records from Anthropic (2026), Google (2026a, 2026b), and OpenAI (2026).

Each lesson was parsed into a structured dataset containing provider, course, section, lesson title, content type, and extracted text where available. An AI-assisted, competency-focused prompt then

 ISSN: 2473-4901
v12 n6753

summarized lessons from all three providers by identifying the knowledge or action being taught, the task or context in which it was applied, and any stated expectations for verification, responsibility, or proficiency. Summaries for title-only lessons remained limited to claims supported by the lesson title and its position within the course.

Inductive coding began with the competency summaries, from which recurring competencies were retained as candidate codes. The authors merged closely overlapping codes, divided broad codes when they described distinct forms of practice, and grouped the resulting codes into higher-order dimensions. They then organized the dimensions into Fundamentals, Professional, and Builder tiers based on learner roles, course organization, and the increasing complexity represented across provider curricula, while keeping academic frameworks from defining the industry dimensions during this stage.

The authors reviewed and revised the AI-assisted summaries, candidate codes, dimension labels, definitions, tier assignments, and provider-support judgments. This review improved conceptual clarity, consolidated overlapping dimensions, and translated competency descriptions into observable forms of knowledge or action. A dimension entered the final framework when all three providers supported it or when two providers supplied clear evidence and the uneven coverage was identified explicitly, while provider-specific topics remained findings rather than shared dimensions.

After finalizing the inductive industry framework, the authors mapped each dimension to the academic AI literacy literature. Hackl et al. (2026) served as the main organizing reference because the study provides a comprehensive review of AI literacy frameworks and offers the AI Literacy Heptagon as a compact curriculum-design vocabulary. Each industry dimension was mapped to one or more Heptagon dimensions and assigned an approximate proficiency level using Unaware, Beginner, Intermediate, and Expert. This second-stage mapping locates industry academy content within the Heptagon without determining the original industry codes or dimensions.

4. RESULTS

The main finding is that the three providers teach AI literacy as a tiered set of skills. The tiers differ across providers, although an overall pattern emerges from the data. All users need basic skills for using AI well and responsibly. Professionals need to apply AI to work processes, build repeatable patterns, and supervise more complex AI behavior. Builders need to design, connect, and test AI systems (see Figure 2).

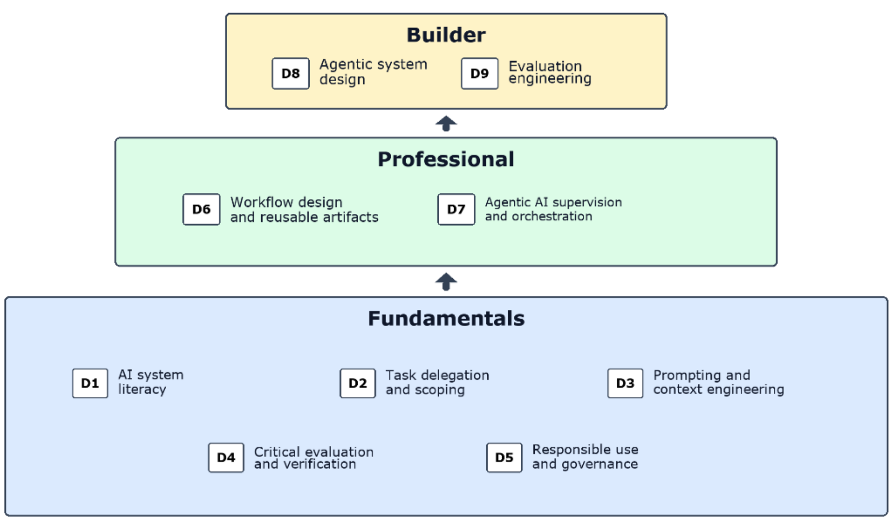

Figure 2. Industry-derived AI literacy framework.

The Fundamentals tier covers five dimensions, beginning with **D1 AI System Literacy**, which means having a working model of how current AI systems behave, including output generation, knowledge and

 ISSN: 2473-4901
v12 n6753

context limits, and predictable failure modes. Anthropic provides substantial coverage through AI Capabilities and Limitations and the AI Fluency Framework Foundations. OpenAI's ChatGPT Fundamentals and Prompting materials and Google's Introduction to AI and Discover the Art of Prompting provide partial coverage, usually as background for effective tool use. **D2 Task Delegation and Scoping** means deciding which work is suitable for AI, defining the goal, choosing the tool, decomposing the work, and retaining human judgment where needed. Anthropic provides substantial coverage through the Delegation pillar. Google's Maximize Productivity with AI Tools provides substantial coverage through task-suitability and human-involvement lessons, while OpenAI's Work Users materials provide partial coverage through use-case identification. **D3 Prompting and Context Engineering** covers clear instructions, examples, documents, constraints, output formats, iteration, and context management. Coverage of D3 is especially even because Anthropic teaches Description and prompting, OpenAI teaches Prompting, Customizing ChatGPT, and Canvas, and Google teaches the full Discover the Art of Prompting course and advanced prompting in AI Fundamentals. **D4 Critical Evaluation and Verification** means checking claims, reasoning, sources, and fitness for use while retaining responsibility for the result. Anthropic's Discernment materials and OpenAI's Responsible Use, Deep Research, and Web Search materials provide substantial coverage. Google provides substantial coverage through human-oversight lessons and source-aware research tasks. **D5 Responsible Use and Governance** covers privacy, security, bias, fairness, disclosure, organizational rules, and human accountability. Anthropic's Diligence materials and OpenAI's Responsible Use of ChatGPT at Work provide substantial coverage. Google's Use AI Responsibly course provides substantial coverage through bias, harm, privacy, and security content.

The Professional tier covers two dimensions, beginning with **D6 Workflow Design and Reusable Artifacts**, which means improving repeated work through reusable prompts, instructions, assistants, templates, knowledge bases, and human-AI process design. Anthropic and OpenAI provide substantial coverage through Agent Skills, Claude Cowork, Projects, Skills, Custom GPTs, and Tasks. Google provides substantial coverage through Gems, NotebookLM knowledge hubs, and AI Studio apps, using different product vocabulary for the same persistent and shareable workflow concept. **D7 Agentic AI Supervision and Orchestration** means defining an agent's task, setting boundaries and permissions, reviewing plans, monitoring multi-step activity, and checking the result. Anthropic's Claude Cowork and agent-supervision materials and OpenAI's Workspace Agents, Tasks, and Deep Research materials provide substantial coverage. Google defines agents in AI Fundamentals, but the selected corpus contains minimal instruction on supervising their actions.

The Builder tier covers two dimensions, beginning with **D8 Agentic System Design**, which means designing agent workflows in code through tools, function calls, routing, handoffs, retrieval, and multi-step plans. Anthropic provides substantial coverage through Building with the Claude API, Model Context Protocol, and Claude Code 101, while OpenAI provides substantial coverage through Building Agents, Codex, and MCP for Builders. Google does not cover this dimension in the analyzed corpus. **D9 Evaluation Engineering** means building test datasets, rubrics, graders, failure analyses, model comparisons, and regression checks. Anthropic's API and cloud-platform materials and OpenAI's Builder Bootcamp, AI App Development, and Model Optimization materials provide substantial coverage. Google also does not cover evaluation engineering in the analyzed corpus.

A cross-provider comparison of the nine industry-derived dimensions appears in Table 1, where each cell pairs the assigned coverage level with the course or module evidence supporting that judgment.

ISSN: 2473-4901
v12 n6753

Table 1. Provider support for the industry-derived AI literacy dimensions.

Dimension	Anthropic Academy	OpenAI Academy	Google AI Academy
D1 AI System Literacy	**Substantial** (AI Capabilities, AI Fluency Foundations)	**Partial** (ChatGPT Fundamentals, Prompting)	**Partial** (Introduction to AI, M2, Art of Prompting, M1)
D2 Task Delegation and Scoping	**Substantial** (AI Fluency Foundations: Delegation)	**Partial** (Work Users: use-case identification)	**Substantial** (Maximize Productivity, M3)
D3 Prompting and Context Engineering	**Substantial** (AI Fluency Foundations: Description, prompting lessons)	**Substantial** (Prompting, Customizing ChatGPT, Canvas)	**Substantial** (Art of Prompting, AI Fundamentals, M4-M5)
D4 Critical Evaluation and Verification	**Substantial** (AI Fluency Foundations: Discernment)	**Substantial** (Responsible Use, Deep Research, Web Search)	**Substantial** (Human oversight, Research and Insights)
D5 Responsible Use and Governance	**Substantial** (AI Fluency Foundations: Diligence)	**Substantial** (Responsible Use of ChatGPT at Work)	**Substantial** (Use AI Responsibly)
D6 Workflow Design and Reusable Artifacts	**Substantial** (Agent Skills, Claude Cowork, Projects and Skills)	**Substantial** (Skills, Custom GPTs, Projects, Tasks)	**Substantial** (Gems, NotebookLM hubs, AI Studio)
D7 Agentic AI Supervision and Orchestration	**Substantial** (Claude Cowork, agent-supervision modules)	**Substantial** (Workspace Agents, Tasks, Deep Research)	**Minimal** (Agent definition, no supervision instruction)
D8 Agentic System Design	**Substantial** (Claude API, MCP, Claude Code 101)	**Substantial** (Building Agents, Codex, MCP for Builders)	**Not covered**
D9 Evaluation Engineering	**Substantial** (Claude API evaluation, Bedrock and Vertex AI)	**Substantial** (Builder Bootcamp evals, AI App Development, Model Optimization)	**Not covered**

When the nine industry dimensions are mapped onto the AI Literacy Heptagon (Hackl et al, 2026), they form a curriculum profile (see Figure 3). The profile is strongest in Application Proficiency, Critical Thinking Ability, and Integration Skills, where the industry academies reach the Expert level. It reaches the Intermediate level in Technical Knowledge and Skills and Ethical Awareness and Reasoning. It reaches the Beginner level in Social Impact Understanding and Legal and Regulatory Knowledge. The mapping is based on the dominant competency implied by each industry dimension and, where relevant, the secondary competencies needed to perform that work.

ISSN: 2473-4901
v12 n6753

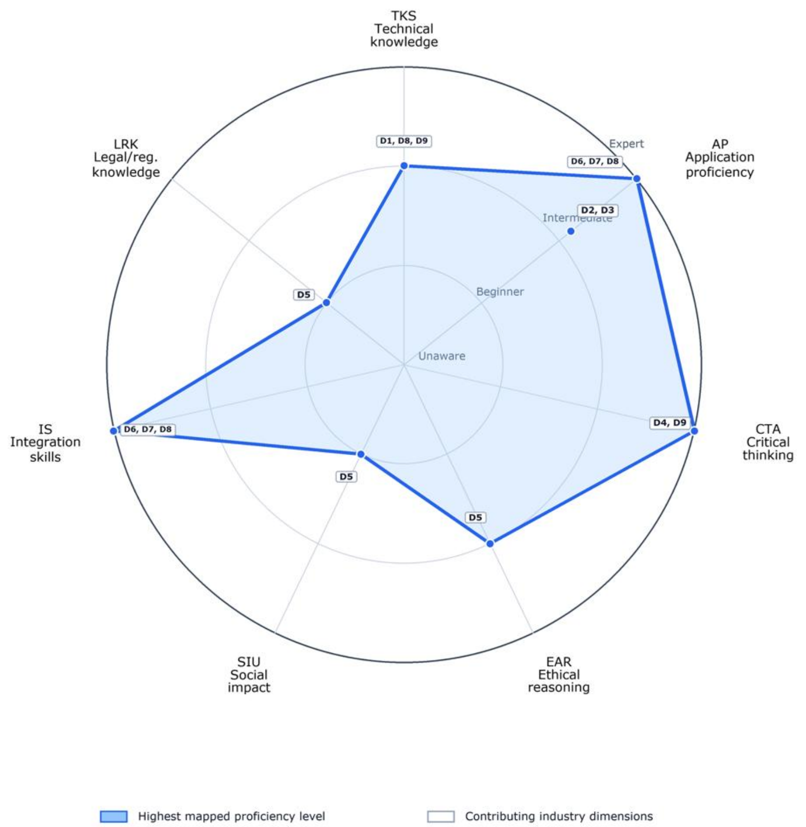

Figure 3. Mapping the industry-derived dimensions to the AI Literacy Heptagon (Hackl et al., 2026)

The mapping shows that each of the first four dimensions (D1–D4) maps directly onto one dimension in the AI Literacy Heptagon, while the remaining five dimensions (D5–D9) span multiple dimensions in the Heptagon. **D1 AI System Literacy** contributes mainly to Technical Knowledge and Skills at the Intermediate level. **D2 Task Delegation and Scoping** and **D3 Prompting and Context Engineering** contribute mainly to Application Proficiency at the Intermediate level. **D4 Critical Evaluation and Verification** maps strongly to Critical Thinking Ability at the Expert level. **D5 Responsible Use and Governance** maps across Ethical Awareness and Reasoning at the Intermediate level and across Social Impact Understanding and Legal and Regulatory Knowledge at the Beginner level. **D6 Workflow Design and Reusable Artifacts** maps to Application Proficiency and Integration Skills at the Expert level. **D7 Agentic AI Supervision and Orchestration** maps to Application Proficiency and Integration Skills at the Expert level. **D8 Agentic System Design** maps to Technical Knowledge and Skills at the Intermediate level and to Integration Skills and Application Proficiency at the Expert level. **D9 Evaluation Engineering** maps to Critical Thinking Ability and Technical Knowledge and Skills at the Expert level.

5. DISCUSSION

The results show how the industry academies define AI literacy from the ground up, with substantial overlap between the industry corpus and academic AI literacy frameworks in three areas. First, both traditions treat critical evaluation as central, with the industry academies teaching users to check AI output before relying on it while the academic synthesis treats critical thinking as a core AI literacy dimension. Second, both traditions treat responsible use as necessary, although academic work often uses the language of ethics, law, and social impact while the academies often use the language of privacy, data handling, disclosure, and user responsibility. Third, both traditions recognize practical use and integration, with the industry corpus giving concrete examples of AI-supported work and the

 ISSN: 2473-4901
v12 n6753

Heptagon giving curriculum designers a way to place these examples within Application Proficiency and Integration Skills.

Because these academies are created by firms that develop and deploy leading AI systems, their course choices also function as signals of industry priorities. The emphasis on prompting, reusable workflows, agent supervision, and evaluation suggests that industry increasingly expects AI competence to extend beyond one-time tool use. Information systems programs can use these signals when updating curricula and defining role-specific proficiency targets, while IS researchers can study how provider-defined skills move into organizational routines, job expectations, governance practices, and system design. These signals require critical interpretation because vendor academies are shaped by product strategies and give less attention to some social, legal, and regulatory dimensions that higher education must retain.

The clearest gap for academic curricula is the lack of detail around current professional practice. The industry academies teach context management, reusable prompt and project artifacts, custom assistants, agent supervision, and evaluation routines as ordinary parts of AI work. These practices belong naturally in the Heptagon, especially under Application Proficiency and Integration Skills, but they need more specific curriculum targets. A higher-education program using the Heptagon could ask whether students can create reusable instructions for a repeated task and manage context across a long AI-supported project. It could also ask whether students can decide when an agent should act, when a human should approve the next step, and whether an AI workflow still works after a model or prompt change.

The Professional tier is especially important for curriculum design because it names a middle layer between basic AI use and software development. Many students will not build production AI systems in code, although they will need to design repeatable AI-supported work, build shared templates and review practices, and supervise agents that can take several steps across tools. Academic curricula can use the Heptagon to locate this work under Integration Skills. The curriculum itself should also make these practices visible and assessable.

The Builder tier also raises questions for higher education because the industry corpus treats agentic system design and evaluation engineering as part of advanced AI literacy. These skills map most clearly to Technical Knowledge and Skills, Application Proficiency, Integration Skills, and Critical Thinking Ability, yet they may not belong in every course or program at the same level. A business student, an education student, and a computer science student may all need to understand that AI systems require testing, even when only some students need to build evaluation pipelines. The Heptagon's proficiency levels are useful here because they let programs choose different target levels rather than treating AI literacy as the same for every student.

The mapping also shows where industry academies could improve, especially in Social Impact Understanding and Legal and Regulatory Knowledge. The academies often focus on what the individual user or organization should do, but they give less attention to broader effects on work, education, institutions, inequality, and public life. Industry courses may include policy reminders, privacy warnings, and enterprise guidance, but they rarely treat law and regulation as part of AI literacy itself. Higher education cannot ignore these areas, and industry academies would be stronger if they taught them more directly.

Technical knowledge is another area where the two lenses differ, since the industry corpus often teaches a practical mental model of AI behavior rather than a deeper view of algorithms, training data, and system design. This practical focus is useful for many users because a person does not need to know every detail of model training to use an AI tool responsibly. Higher education still has a broader duty to help students understand how AI tools work, why they fail, and how data, design, and deployment choices shape outcomes.

The comparison suggests that academic and industry AI literacy answer different questions. Academic work asks what a person in higher education should understand, evaluate, and act on when using AI. The industry academies ask what a person can do with AI tools and how that person can remain responsible for the result. Each view remains incomplete without the other. A user who can build

ISSN: 2473-4901
v12 n6753

workflows and supervise agents without understanding social harm, law, bias, or broader effects has a narrow form of skill. A user who understands AI in social and technical terms may still struggle to use the tools effectively without prompting, context management, and output review skills.

This study has several limits, beginning with the fact that the academy corpus is a snapshot from May 2026 and that the courses change often, so the findings describe that period rather than a fixed curriculum. In addition, many Anthropic and OpenAI developer lessons were video-heavy and available only through titles and course structure, which likely makes the Builder tier less complete than the Fundamentals and Professional tiers. The framework was developed with AI-assisted summarization followed by author review, but lessons were not independently double-coded and inter-rater reliability was not calculated. Future work should therefore add independent human coding of a lesson sample and compare coder agreement. Finally, the academic mapping is organized through a recent synthesis of the literature, which avoids repeating separate reviews of older frameworks but also means that details from individual older frameworks are discussed mainly through their place in that synthesis.

6. CONCLUSION

This paper used an inductive analysis of Anthropic Academy, OpenAI Academy, and Google AI Academy, conducted in May 2026, to define AI literacy as taught by major AI providers, resulting in a three-tier, nine-dimension framework organized around Fundamentals, Professional, and Builder levels. The mapping to academic AI literacy work shows strong agreement on the broad foundations of this framework, especially critical evaluation, responsible use, legal and ethical reasoning, and workflow integration. The industry academies add more detail about present-day practice, including context management, reusable AI artifacts, agent supervision, and evaluation engineering. The main implication is that higher education should combine the broad educational model developed in the literature with the more concrete skills now being taught by AI providers. Students need to understand AI, judge it carefully, use it responsibly, and work with it effectively.

7. REFERENCES

Almatrafi, O., Johri, A., & Lee, H. (2024). A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts. *Computers and Education Open*, *6*, 100173. https://doi.org/10.1016/j.caeo.2024.100173

Anthropic. (2026). *Anthropic Academy* [Online learning platform]. Retrieved May 15, 2026, from https://www.anthropic.com/learn

Google. (2026a). *Google AI Essentials Specialization* [Online course]. Coursera. Retrieved May 15, 2026, from https://www.coursera.org/specializations/ai-essentials-google

Google. (2026b). *Google AI Professional Certificate* [Online course]. Coursera. Retrieved May 15, 2026, from https://www.coursera.org/professional-certificates/google-ai

Hackl, V., Mueller, A., & Sailer, M. (2026). The AI literacy heptagon: A structured approach to AI literacy in higher education. *Computers and Education: Artificial Intelligence*, *10*, 100540. https://doi.org/10.1016/j.caeai.2026.100540

Kassorla, M., Georgieva, M., & Papini, A. (2024). *AI literacy in teaching and learning: A Durable Framework for Higher Education*. EDUCAUSE. https://www.educause.edu/content/2024/ai-literacy-in-teaching-and-learning/defining-ai-literacy-for-higher-education

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-16). ACM. https://doi.org/10.1145/3313831.3376727

Miao, F., & Cukurova, M. (2024). *AI competency framework for teachers*. UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000391104

ISSN: 2473-4901
v12 n6753

Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, *2*, 100041. https://doi.org/10.1016/j.caeai.2021.100041

OECD & European Commission. (2025). *Empowering learners for the age of AI: An AI Literacy Framework for primary and secondary Education.* https://ailiteracyframework.org/wp-content/uploads/2025/05/AILitFramework_ReviewDraft.pdf

OpenAI. (2026). *OpenAI Academy* [Online learning platform]. Retrieved May 15, 2026, from https://academy.openai.com/