Defining an AI-Literacy Course for Dual-education Programs, in Computer/data Science: A Case Study

Generative AI, especially Large Language Models, has the potential to transform learning and evaluation processes. However, institutions often focus primarily on regulations to guarantee academic integrity rather than on enabling effective use. Generic AI-Literacy programs are not sufficient, adoption and use of generative AI are wildly different between domains. As a literature review shows, use of AI should also be taught with realistic, domain specific hands on exercises. For Software Engineering students it is not sufficient to teach AI-Literacy at the end of a (bachelor) program but students need to be able to use AI in the right way – compliant and effective – from the beginning on. We present a three-day AI literacy curriculum developed at DHBW (Duale Hochschule Baden-Württemberg) covering technical, practical, and legal aspects. The first part introduces AI fundamentals, including neural network and transformer architectures, as well as applications such as RAG systems and agentic AI. The second part is domain-specific, addressing practical applications like AI-assisted scientific writing or code generation. The third part covers ethical and regulatory topics including copyright, examination law, and university regulations, concluding with a community-based process for developing a code of conduct. DHBW's dual-education model, where students alternate between university and industry partners, presents unique challenges: students bring extensive real-world experience using generative AI for tasks like software development and presentation preparation. This expertise makes predefined rules difficult to enforce and renders fixed curricula impractical. We therefore adopted an inclusive, community-based approach to developing guidelines—an approach that has proven effective in gaining acceptance. By examining benefits and challenges, including academic integrity and accessibility, this paper provides a base for discussions about AI integration in higher education.

Keywords: *Generative AI, Dual Education, Academic Integrity, AI-Literacy*

Introduction

"If gen AI has a killer application, it's software development—one of the most profound shifts in the history of programming." (McKinsey, 2026)

The use of and research about generative AI is one of the hottest topics in computer science, according to the Stanford AI index the share of AI-related publication is approaching 50% of total computer science publications (Stanford University, 2026). This development is not just about tool use, but a fundamental change in how software is developed: we are moving towards a "symbiotic partnership between human software developers and AI" in this area (Terragni, 2025). Generative AI (genAI) is relevant in other sectors too, but we will focus on this domain, because we must change the way we teach these topics.

The traditional view in Computer Science (CS) education, that the details of how AI works (and how to use it) should be taught maybe in the third year or later is not adequate today, especially in a dual study program like those at Duale Hochschule Baden Württemberg (DHBW). DHBW with its 33.000 students and more than 230.000 Alumni is one of the largest universities in Germany and the largest university with an integrated dual study program worldwide (DHBW, 2026).

Since our students spend only half of their time with us and the other half in their partner company, working as software engineers resp. data scientists in the study programs of the authors, our students use these techniques and generate code and related artifacts on a regular base. Because of that, we cannot just tell them, not to use AI in university; students won't accept that based on their work experience. They will use AI from the beginning on, so we have to teach them how to use it correctly in the first year. The problem with that is, that there is a big gap in knowledge between the students, which can not be closed just by short introductions or demos (Ma, 2026).

The questions answered in this paper are: How can we prepare our new software engineering students to use AI in a compliant and meaningful way?

Do we need special content or can we use a generic, domain agnostic approach? How can we include their industry experience?

Domain specific AI literacy is underexplored in current research (Knoth,2024, p.4), we contribute by improving that situation for teaching computer science, especially programming: we present a case study in how AI-Literacy can be taught in a dual study program to include the diverse experiences, the students get in their respective companies. This allows to include very different perspectives in an inclusive way.

Beginning with the introduction we will discuss current developments in this area and motivate our work by discussing relevant literature. Then we present and discuss the structure of our course and the process to include student experience.

Related work on AI literacy programs

There is a body of relevant literature about AI-literacy education starting with the definition by (Long and Magerko, 2020), who introduced AI literacy as a set of competencies that allow people to communicate, work and live with AI, one of them being able to identify problems where AI excels and use it on these. Building on this work, several competency frameworks have been proposed like the European DigComp 2.2 framework (Vuorikari et al., 2022) and the EU-AI-Literacy framework (EU-AILIT, 2025), generative-AI competencies of Annapureddy et al. (2025), and the holistic assessment matrix of Knoth et al. (2024) that combines generic, domain-specific and ethical dimensions. Faruqe, Watkins and Medsker (2021) argue that a usable AI-literacy framework must go beyond conceptual catalogues and be turned into a multi-level competency model with assessable behavioural anchors.

At the curriculum level, programs vary in scope. On the short-format end, probably all universities run one-day or one-week intensives such as the one-day

AI-literacy workshop offered by UC Davis (UC Davis, 2026) or the four-hour scalable “AI Literacy for All” course of the Digital Education Council (DEC, 2026). Pre-arrival online short courses such as those of Durham ISC (Durham, 2026) follow the same logic of providing a minimal shared baseline before students start their studies. At the other end of the spectrum, Southworth et al. (2023) describe the University of Florida’s “AI Across the Curriculum” initiative as a multi-year integration across all 16 colleges, and dedicated full-semester courses such as Yale’s “AI for Future Presidents” (Candon et al., 2025) or the “Essentials of AI for Life and Society” at UT Austin (Biswas et al., 2025) deliver AI literacy to general audiences without programming prerequisites. Across this range, recent reviews of higher-education offerings consistently report that approaches combining conceptual knowledge with hands-on, experiential learning are more effective than either alone see for example (Hong, 2025) — a finding that motivates the workshop-heavy design of Pillar 2 in our own course.

For computer-science and data-science students the situation is qualitatively different. AI courses for CS majors have long existed, but they typically emphasize algorithms and mathematical foundations and treat critical and ethical questions as an optional add-on at the end of the semester. Recent work argues that this is no longer sufficient once students routinely use generative AI in their daily programming work. Ma, Koedinger and Wu (2026) show in a controlled study with programmatic data-science tasks that lightweight demonstrations of LLM use improve surface fluency but are insufficient to develop “resilient AI use skills”: students vary in how they form intent, decompose prompts, interpret output and validate results, and prior technical experience remains a significant predictor of success. They conclude that AI literacy for CS and data-science students needs explicit scaffolding around exactly these stages — a conclusion that directly relates to the prompt-engineering and peer-review structure of our Pillar 2.

Kennedy and Gupta (2025) reach a similar conclusion at curriculum level with their AI & Data Acumen Learning Outcomes Framework, which scaffolds AI competencies across proficiency levels and knowledge dimensions and explicitly balances technical skills with ethical and sociocultural awareness. All of them are broad, with a wide societal perspective but don’t give enough domain specific advice, how to prepare students to support their work using genAI.

AI literacy in vocational and work-integrated settings is still underexplored. Hong (2025) develops and validates a competency-based ladder pathway for higher vocational students, explicitly aligned to industry competency requirements; the study reports significant performance gaps between industry expectations and current educational outcomes, particularly in technical skills, critical thinking and ethical awareness. Our course structure focuses on these topics. While this study is based on a large number of subjects and provides a complex, powerful pathway to teach AI literacy, it is not domain specific and is more suited to holistic multi-year study programs than to an introductory course in the first year. But there are important and relevant results, like that experiential learning through hands-on labs is the most effective and most (student) satisfying method evaluated – so we do that too.

The work of Hong et al. gives an interesting path to explore in the long term, but it doesn't solve our problem: we are not aware of any published curriculum that addresses AI literacy specifically for dual-study CS and data-science students practice integrated study programs, where the cohort already brings substantial industry experience into the first year. The case study presented in this paper is intended to close that gap.

Design of the course structure

Why did we develop a new course structure? Our official DHBW position on AI use (DHBW, 2024) for example describes the need to integrate generative AI in our programs and also addresses the need to do that with respect to the needs of our dual partner companies. But it only expresses that we want to use and teach AI by integrating it in the curriculum, not how to do that. Those guidelines are far too abstract to be implemented directly.

That is the case not only in our university, but also in other work. For example, the EU-AI Literacy Framework states, that the General „Critical-Thinking“ skill is important in evaluating AI-Output (EU-AILIT, p. 21) but it is not clear, how this skill should be developed and to what kind resp. in what depth it should be applied. To be used effectively, all students should do that to the same rules.

Because of our students diverse experience from industry, acceptance of predefined rules is low, at least in our experience. To make things even more complicated, acceptance and use of gen AI differ wildly between domains (Knoth, 2024, p4). This is the reason, why we used a collaborative approach to define guidelines for acceptable use and the way to document that, see section 4 (**Code of Conduct**).

The EU AI-Literacy Framework defines „AI-Literacy is an educational Priority“ (EU-AILIT, 2025, p. 8) and focuses on ethical and societal aspects. Our work focuses on creating artifacts like program code, homework assignments or thesis papers with AI. This is covered in the Domain "Creating with AI“ (EU-AILIT, 2025, p. 30ff). For a discussion of the importance of AI literacy to the general public see for example (Long, 2020).

It is important to "Collaborate with AI to create and refine original ideas while considering issues of ownership, attribution, and responsible use." (EU AILIT, 2025, p. 21) but we must define, what usage is acceptable from an academic point of view. The framework states an important point, that we use as a guideline for our students: "By engaging creatively and responsibly with AI systems, learners stay accountable for the ideas they shape and share." Students have to take ownership of the results of their interaction with AI. That allows the students to define for themselves, how they get and hold this ownership in every individual case.

If AI is used at will by students when creating artifacts, the results get worse (without the students noticing!) compared to regulated use: "Self-regulated students trained less, reported a lower sense of accomplishment, and became increasingly reliant on AI, despite being aware of its potential harms.” (Poulidis, 2025). Therefore it is important to practice the use of AI with the students in a controlled way to show then meaningful use, that keep the learning effect while making the process more effective.

A (if not “the”) important point in teaching AI use is to make students realize, that the use of genAI could hinder the learning process. Shaw and Nave (Shaw, 2026) measure “cognitive surrender in their study. They extend the notion of the fast and the slow thinking system with a third: external artificial intelligence. In studies they analyze, what happens, when people use AI in that way resulting in "cognitive surrender-adopting AI outputs with minimal scrutiny, overriding intuition (System 1) and deliberation (System 2). [...] Engaging System 3 also increased confidence, even following errors.” (Shaw, 2026).

When AI produces correct results, the result of the creative process is better than the one obtained without AI. But, if the result is not correct, people tend to believe it despite being wrong (Dell'Acqua, 2023). The better the results from AI get, the more users tend to believe them, which leads to superficial reviews and undetected errors: the better AI gets - and it is capable these days - the more errors will go undetected. Students must learn – by experience – that review is necessary but doesn’t guarantee correct results.

To reduce this effect, it is critical, that students learn, how to prompt in a way, that reduces typical AI-errors (NIST, 2024) like hallucination and over-generalization. These errors are just characteristics of how Large Language Models (LLMs) are trained and answers are generated, so students need to understand the process a LLM works in enough detail to create prompts, that are specific enough to prevent those problems from happening. And be prepared to review answers for their presence.

In developing the structure for our AI-Literacy course, we followed the central idea of Knoth et al: it is not sufficient to teach generic AI literacy and ethics, but it is necessary to include “domain specific [...] AI competencies tailored to the needs and applications within specific professional domains” (Knoth, 2024). In that paper an extensive discussion about the importance of domain specific AI literacy can be found.

Not only this domain specific literacy is important, but “the applications of AI and the competencies required for each may vary wildly across domains and disciplines making different aspects of AI literacy differently significant depending on the professional domain” (Knoth, 2024, p4).

So, it is not sufficient to create one AI literacy course or curriculum, different courses are needed for different domains.

In addition, the adoption of AI differs massively between domains, software engineering probably being at the forefront of adoption. According to (Aikido University, 2026) 85-92% of software engineers use AI regularly to create code and other artifacts. According to (Saran, 2026) at least 50% of new code at alphabet, Google’s parent company, is written by coding agents. According to (Novet, 2025) Satya Nadella, the CEO of Microsoft, said that 30% of Microsoft code is generated too, Meta’s Marc Zuckerberg said “Our bet is sort of that in the next year probably … maybe half the development is going to be done by AI, as opposed to people, and then that will just kind of increase from there,”. So, we are in urgent need to teach students, how to do that effective and efficient.

This leads to our AI literacy course, consisting of three pillars:

Table 1. *Structure of the AI-literacy course — three pillars across a three to four-day block, with two parallel tracks in Pillar 2*

Pillar	Day	Theme	Scope	Content / Key topics
1	1 (0.5)	**Foundations of AI** *How does it work?*	Generic *(discipline-independent)*	• Fundamental concepts and types of AI • Distinction between classical search, narrow AI and AGI • Evolution from GOFAI to modern generative AI • Machine learning and neural network basics • How large language models process and generate text
2	2 or 3	**What's in it for me?** *Practical applications*	Domain-specific *either* computer science / data science *or* business administration	**Common to both tracks** • Hands-on workshops with real-world tasks from university or the partner companies • Effective prompt engineering for the respective domain • Peer review of AI-assisted artefacts **Computer science / data science track** • AI-assisted code generation, refactoring and optimization • AI in SE process (requirements engineering to

Pillar	Day	Theme	Scope	Content / Key topics
				test and deployment) • Use of AI in data-science workflows (exploration, cleaning, modelling) • Product ideation and pitching of technical solutions **Business administration track (alternative)** • AI-assisted business development and ideation • AI-assisted drafting of customer reports and proposals • Business-model sketching and pitching with AI
3	4 (0.5)	**What is difficult?** *Critical perspectives*	Generic *(ethics and regulation)*	• Ethical and legal frameworks for AI (EU AI Act, copyright, examination law) • Recognising and addressing bias in models and training data • Privacy, data protection and transparency • Typical AI failure modes (hallucination, over-generalisation, cognitive surrender) • Community-based development of a local code of conduct

Pillar	Day	Theme	Scope	Content / Key topics
				• Group discussions and case studies on responsible use

The design follows the central observation of Knoth et al. (2024) that a holistic AI-literacy programme must combine *generic*, *domain-specific* and *ethical* competencies, and that the relative weight of these dimensions varies considerably between disciplines[1]. The course is delivered as a three to four day block and is positioned in the first year of our software-engineering and data-science programs, so that students enter their next industry phase with a shared baseline for working with generative AI. The extension to Business Administration programs is obvious: instead of programming, topics like business model generation or product development can be covered. For external (international) students, the program can be extended with a company visit on day one, which is not necessary for our students – they have extensive experience from industry. So, the total duration is 3 days for internal students and 4 days for external students. Combined with a one day workshop on a real company problem, this can be packaged as a 3 ECTS Blended Intensive Program (Erasmus+) or a Micro credential.

Pillar 1 (Day 1) establishes a common technical baseline. Although our students arrive with substantial practical AI experience from their partner companies, that experience might be uneven and based on folk-models of how large language models behave. This pillar therefore covers the evolution from GOFAI to modern generative systems, the basics of machine learning and neural networks, and — in enough depth to inform later prompting decisions — how transformer-based language models actually generate text. This common ground is necessary for understanding why typical failure modes such as hallucination and over-generalization occur and where they are most likely to appear (NIST, 2024).

Pillar 2 (Days 2–3) deepens the understanding of the internals of LLMs and is the domain-specific core of the course and the part most heavily shaped by the dual-education context. It is offered in two parallel/alternative tracks. The computer-science / data-science track is built around AI-assisted code generation and review as well as the use of AI in the whole software engineering process (like test or requirement engineering, depending on the students background) and/or the use of AI in typical data-science workflows; software engineering and data science are at the forefront of AI adoption, and a substantial share of code at major industry players is already AI-generated.

In a parallel track, scientific writing using AI is covered. Both tracks share a common scaffold of hands-on workshops with realistic tasks drawn, where possible, from the students' partner companies, prompt engineering tailored to the domain, and systematic peer review of AI-assisted artefacts. The peer-review

[1]Although Knoth et al focuses on AI literacy assessment and not on teaching it, the argument about the importance of domain specific parts are valid here too.

element is used to counter the “cognitive surrender” effect described in (Shaw, 2026), under which uncritical acceptance of AI output produces measurably worse results.

Pillar 3 (Day 4) returns to a generic perspective, but with the practical experience of Days 2–3 as an anchor. It covers topics like the legal and ethical framework (EU AI Act, copyright and examination law, university regulations), bias and transparency, and the typical risks of generative AI use in academic and professional settings. That leads to a community-based exercise in which the cohort co-develops a local code of conduct for AI use in their study program (see the next section for details). In our experience this inclusive format produces higher acceptance of the resulting rules than the top-down guidelines that are typical at this level.

The combination of generic foundations, domain-specific practice and generic ethical reflection follows the holistic-literacy logic of Knoth et al. (2024) and goes one step further by each pillar reinforcing the next: the foundations make the practical labs more critical, the labs make the ethical discussion more concrete, and the co-developed code of conduct carries the entire course back into the industry phase.

Code of Conduct

Since we (at least in our university as of May 2026) don’t have an agreed upon definition of acceptable use of AI in university (besides very general rules from examination regulations about academic integrity) on the one hand and our students having a growing body of practical experience using generative AI to do their work in their partner company on the other, we use an individual (per cohort of some 30 students) code of conduct that we develop together with the students in an iterative process[2].

We start that part (the third pillar of our AI Literacy course) by briefly introducing the legal and regulatory (general and university specific) aspects.

Then we explore the range of possibly meaningful uses of generative AI by letting the students (in small groups like 2-4 persons) define tasks that might be supported by genAI. This results typically with some 30 ideas (see Figure 1). We ask the students to position those ideas on a board spanning the range from perfectly OK (left) to absolutely forbidden (right). Using a virtual board (miro in this case) lowers the barrier to show things, the students are not sure about, so we prefer that. Figure 1 visualizes the result of the following process. In the beginning most of the ideas group as a colored blob in the middle, probably because the students sort out obvious candidates for “good” or “bad” by themselves.

[2]One might argue, that having different agreements for different groups of students is asking for trouble regarding fairness of use. But in spite of having an absolute truth to compare against we consider this as different realities which exist in parallel and will be interesting to analyze for temporal development in the future.

Figure 1. *Miro board used to develop the code of conduct (symbolic)*

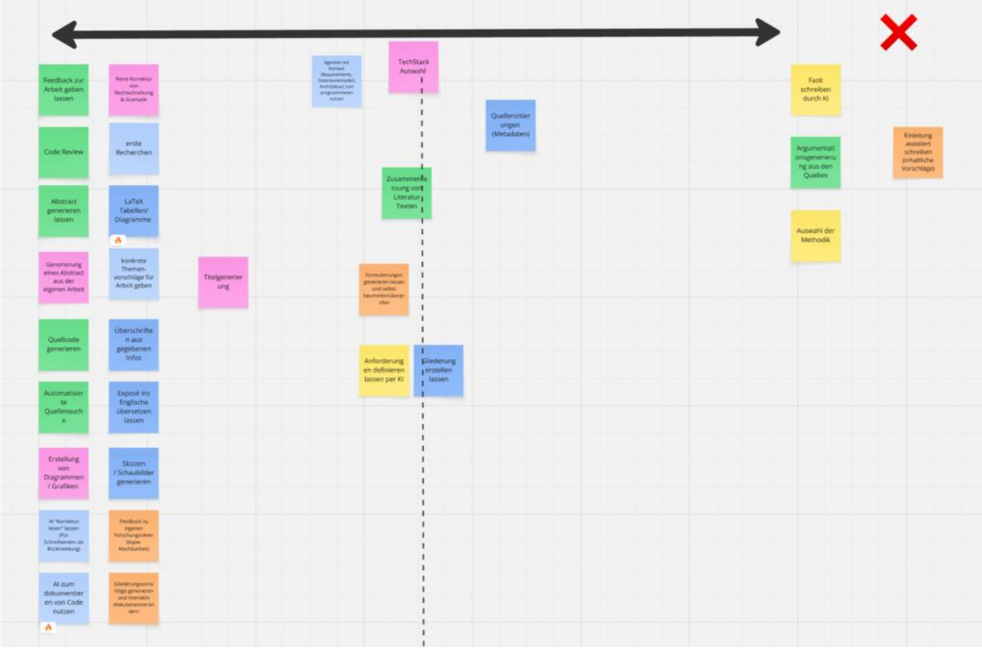

We start the process by asking the students to identify ideas, which they are pretty sure, that they are “forbidden” or “good”, discuss them individually and move them around accordingly. After a while of discussion (typically maybe 30 minutes) there are some (typically approximately 30%) of the ideas left in the middle. These are the interesting ones ….

This leaves questions about the goodness of tasks like summarizing literature, create metadata for references, select a tech stack, define requirements for their work etc. These questions are hard to answer in a general way. If these uses are acceptable or not depends on the domain or the subject of their work if it’s a thesis paper. If it is about an assignment in the context of a lecture it depends on what students should learn/show in this assignment. In any case, the use in these cases must be documented appropriately.

This is the area, where we must discuss, how to come to a decision, if AI use is good vs. acceptable vs. bad. What we want to teach the students is to reflect on their use of genAI to decide in these cases and when and how to discuss that with their coaches/supervisors. The result of this process is converted to a document that can be used as a reference for the respective assignment/thesis.

We ran this process for four times now (May 2026) and it produced very similar outcomes. We were able to reach consensus on what is OK and what not in all three cohorts of computer science students with whom we tried this approach. Of course, this is not a final result but a snapshot of work in progress.

Conclusions

Of course, our 3-4 day course is not sufficient to enable computer science (or other) students to use generative AI to the full extent possible. It is meant to give them a head start at the beginning of their studies in the first year. Since generative AI has implications on and connections to nearly everything taught in the different programs, a much broader approach has to be deployed "The ultimate goal of AI across the Curriculum is the creation of an AI-ready workforce covering the essential 21st-century competencies identified as workforce and government needs worldwide." (Southworth, 2023).

But especially in computer science and especially in our dual study programs, we cannot wait for that to happen. That is the reason, why we developed the course described in this paper. Its structure is flexible enough to be adapted to other programs like data science or Business Administration.

Maybe an incremental approach - including the expertise and experience, not only of the teachers, but also of the students - to the development of AI Literacy skills is more appropriate considering the rapid development in this area, especially in software engineering. At least in our experience and in a dual setting.

References

Annapureddy, Ravinithesh, Fornaroli, Alessandro, Gatica-Perez, Daniel, *Generative AI Literacy: Twelve Defining Competencies*, Digital Government: Research and Practice, Volume 6, Issue 1, Article 13, 2025, available online https://doi.org/10.1145/3685680

Biswas, Joydeep, Fussell, Don, Stone, Peter, Patterson, Kristin, Procko, Kristen, Sabatini, Lea, Xu, Zifan, *The Essentials of AI for Life and Society: An AI Literacy Course for the University Community*, Proceedings of the AAAI Conference on Artificial Intelligence, Volume 39, Issue 28, pp. 28973–28978, 2025, available online https://doi.org/10.48550/arXiv.2501.07392

Candon, Kate, Georgiou, Nicholas C., Ramnauth, Rebecca, Cheung, Jaden, Fincke, Emma C., Scassellati, Brian, *Artificial Intelligence for Future Presidents: Teaching AI Literacy to Everyone*, Proceedings of the AAAI Conference on Artificial Intelligence, Volume 39, Issue 28, pp. 28988–28995, 2025, available online https://doi.org/10.1609/aaai.v39i28.35168

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.* Harvard Business School Working Paper No. 24-013.

DHBW, *Positionspapier zu künstlicher Intelligenz*, Duale Hochschule Baden-Württemberg, 2024, available online https://www.dhbw.de/fileadmin/user_upload/Dokumente/Positionspapiere_und_Strategie/Positionspapier_zu_kuenstlicher_Intelligenz.pdf

DHBW, *Duale Hochschule Baden-Württemberg*, 2026, available online http://www.dhbw.de

Digital Educational Council (DEC), "AI Literacy for All", 2026 https://www.digitaleducationcouncil.com/digital-education-council-ai-literacy-for-all-course

Durham University, "AI Literacy short course ", 2026, https://www.durhamisc.com/programmes/pre-courses/ai-digital-literacy

EU-AILIT, European Commission, An AI Literacy Framework for Primary and Secondary Education, (Review draft), 2025, https://ailiteracyframework.org/wp-content/uploads/2025/05/AILitFramework_ReviewDraft.pdf

Faruqe, Farhana, Watkins, Ryan, Medsker, Larry, *Competency Model Approach to AI Literacy: Research-based Path from Initial Framework to Model*, arXiv preprint arXiv:2108.05809, 2021, available online https://arxiv.org/abs/2108.05809

Hong, Litian, *Development and validation of a competency-based ladder pathway for AI literacy enhancement among higher vocational students*, Scientific Reports, 2025, available online https://doi.org/10.1038/s41598-025-15202-6

Kennedy, Kathleen, Gupta, Anuj, *AI & Data Competencies: Scaffolding holistic AI literacy in Higher Education*, arXiv preprint arXiv:2510.24783, 2025, available online https://arxiv.org/abs/2510.24783

Knoth, Nils, Decker, Marie, Laupichler, Matthias Carl, Pinski, Marc, Buchholtz, Nils, Bata, Katharina, Schultz, Ben, *Developing a holistic AI literacy assessment matrix – Bridging generic, domain-specific, and ethical competencies*, Computers and Education Open, Volume 6, Article 100177, 2024, available online https://doi.org/10.1016/j.caeo.2024.100177

Long, Duri, Magerko, Brian, *What is AI Literacy? Competencies and Design Considerations*, Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), Association for Computing Machinery, 2020, available online https://doi.org/10.1145/3313831.3376727

Ma, Qianou, Koedinger, Kenneth, Wu, Tongshuang, *Not Everyone Wins with LLMs: Behavioral Patterns and Pedagogical Implications for AI Literacy in Programmatic Data Science*, arXiv preprint arXiv:2509.21890, 2026, available online https://arxiv.org/abs/2509.21890

McKinsey & Company, *The AI revolution in software development*, 2026

National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1, NIST Trustworthy and Responsible AI, 2024, available online https://doi.org/10.6028/NIST.AI.600-1

Novet, Jordan, *Satya Nadella says as much as 30% of Microsoft code is written by AI*, CNBC, 2025, available online https://www.cnbc.com/2025/04/29/satya-nadella-says-as-much-as-30percent-of-microsoft-code-is-written-by-ai.html

Poulidis, Stefanos, Bastani, Hamsa, Bastani, Osbert, *Self-Regulated AI Use Hinders Long-Term Learning*, The Wharton School Research Paper, 2025, available online https://ssrn.com/abstract=5604932

Saran, Cliff, *Half of Google's software development now AI-generated*, Computer Weekly, 2026, available online https://www.computerweekly.com/news/366638839/Half-of-Googles-software-development-now-AI-generated

Shaw, Steven D., Nave, Gideon, *Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender*, The Wharton School Research Paper, 2026, available online https://doi.org/10.31234/osf.io/yk25n_v1

Southworth, Jane, Migliaccio, Kati, Glover, Joe, Glover, Ja'Net, Reed, David, McCarty, Christopher, Brendemuhl, Joel, Thomas, Aaron, *Developing a model for AI Across the curriculum: Transforming the higher education landscape via innovation in AI literacy*, Computers and Education: Artificial Intelligence, Volume 4, Article 100127, 2023, available online https://doi.org/10.1016/j.caeai.2023.100127

Stanford University, *Stanford AI Index Report 2026*, Stanford Institute for Human-Centered AI (HAI), 2026, available online https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf

State of AI in Security & Development, *Aikido Security*, 2026, available online https://www.aikido.dev/state-of-ai-security-development-2026

Terragni, Valerio, Vella, Annie, Roop, Partha, Blincoe, Kelly, *The Future of AI-Driven Software Engineering*, ACM Transactions on Software Engineering and Methodology, Volume 34, Issue 5, Article 120, pp. 1–20, 2025, available online https://doi.org/10.1145/3715003

UC Davis, "AI-Literacy and Logic Workshop",2026, https://research.ucdavis.edu/university-writing-program-wins-1-5-million-grant-in-ai-challenge-with-support-from-the-office-of-research-proposal-development-services/

Vuorikari, Riina, Kluzer, Stefano, Punie, Yves, *DigComp 2.2: The Digital Competence Framework for Citizens — With new examples of knowledge, skills and attitudes*, Publications Office of the European Union, JRC128415, 2022, available online https://doi.org/10.2760/115376