Bachelor's Thesis Computing Science

Radboud University Nijmegen

Education following the AI Boom: should the usage of GenAI in education be integrated or banned according to Computing Science experts at Radboud University?

Author:
Mats Simon Wapstra
s1036786

First supervisor/assessor:
dr. I. Wilmont (Ilona)

Second assessor:
dr. B.E. van Gastel (Bernard)

June 16, 2026

Abstract

This thesis examines the impact of large language model (LLM) use on education and AI governance and the consequences for integration of AI within the Radboud University Bachelor of Computing Science. To answer this question, two methods were combined: a rapid review of recent literature on generative AI in higher education and computing science education and a deductive thematic analysis of semi-structured interviews with participants connected to the Bachelor of Computing Science.

The rapid review showed that current literature presents generative AI as both promising and disruptive. Reported benefits include support for learning as a form of online teaching, efficiency, accessibility and student motivation by making faster progress. At the same time, the most common concerns relate to academic integrity, overreliance, unequal access and less meaningful learning processes. The main position that was taken, was neither full rejection nor full adoption, but careful and selective integration supported by policy, increased AI literacy among both students and teachers and assessment redesign.

The interview findings showed a similar pattern. Participants generally described generative AI as too influential to ignore, but also as risky when used without critical judgement or when it replaces important parts of the learning process. Participants strongly emphasized, especially in relation to the Bachelor of Computing Science at Radboud University, the importance of foundational programming knowledge, computational and critical thinking and clear quality awareness before AI can be used responsibly. AI was therefore not seen as a replacement for core learning, but at most as a tool that may become useful once students have first developed the relevant skills themselves.

The combined findings suggest that LLM use already has a significant impact on education and AI governance and that action is needed in the relatively short term. However, the results do not support either full prohibition or broad integration of AI in the Radboud Computing Science curriculum. They point more towards a carefully critical and selective approach. AI would then only be integrated where this can be justified in relation to learning goals and the protection of the core skills as defined by the intended learning outcomes of the Bachelor of Computing Science.

Contents

1 Introduction **3**

2 Related Work **5**
2.1 Educational evolution 5
2.2 Vibe coding 5
2.3 Impact of AI on understanding and learning goals 6
2.4 Generative AI in education: effects on learning and study practices 7
2.5 Generative AI specifically in IT/CS education: effects on programming, assessment, and skill development 8

3 Research **9**
3.1 Interview study 9
3.1.1 Participants 9
3.1.2 Procedure 10
3.1.3 Ethical considerations 11
3.1.4 Analysis 12
3.2 Rapid Review 12
3.2.1 Review question and objectives 13
3.2.2 Eligibility criteria 13
3.2.3 Information sources 14
3.2.4 Databases and search strategy 14
3.2.5 Search queries 15
3.2.6 Procedure 15
3.2.7 Analysis 15
3.3 Trustworthiness 17
3.3.1 Credibility 17
3.3.2 Transferability 17
3.3.3 Dependability 18
3.3.4 Confirmability 18

4 Results **19**
4.1 Interviews 19

4.1.1 Findings 19
4.2 Rapid Review 24
4.2.1 Findings 24

5 Discussion **29**
5.1 Strengths and limitations 29
5.2 Recommendations 30

6 Conclusions **31**

Chapter 1

Introduction

Artificial intelligence (AI) tools have moved from niche research projects to everyday "in-the-pocket" assistants in a remarkably short time[30]. Systems such as large language models can now generate answers, texts and produce working code in seconds. For many people, these tools are no longer an occasional novelty, but a part of their daily life. Students often use them to help with assignments and learning[22]. They are used to brainstorm, debug, summarize and sometimes even complete tasks. This rapid adoption is currently redefining how people learn, how teachers design assignments and how universities try to uphold academic standards.

This tension is clearly visible in computing science education. In the program of the Bachelor of Computing Science at Radboud University, learning outcomes are strongly tied to practicing problem solving through coding and mathematical understanding. Students not only learn syntax and programming patterns, but also develop computational thinking, debugging skills and an understanding of underlying concepts. At the same time, AI tools can already write basic programs, translate pseudocode into code, suggest solutions to standard assignments and instantly explain error messages. The solutions AI provides are currently not always up to code; however, it raises a serious question. When a tool can generate code on demand, what exactly should students still learn to do manually and how should that learning be assessed?

Educators face an uncomfortable trade-off. On the one hand, integrating AI into education could better prepare students for the reality of modern software development, where AI usage is becoming more and more common[18], which is also reflected in the new term "Vibe Coding"[14]. Using AI well, might function as a helpful partner or a code-review assistant that gives feedback and supports independent learning. On the other hand, unrestricted use of AI can undermine the purpose of introductory coding assignments, make it harder to evaluate individual competence and increase inequalities between students who know how to prompt and verify outputs

effectively and those who do not. The central debate is no longer whether AI will be present in the educational ecosystem, but whether universities should deliberately integrate it with clear rules or restrict/ban it to protect core learning goals and fairness.

This thesis investigates this problem from the perspective of experts in Computing Science at Radboud University by asking:

"What is the impact of usage of LLMs on education and AI governance, and what are the consequences for integration of AI usage in the Radboud computing science curriculum?"

Rather than assuming AI is purely beneficial or purely harmful, this thesis focuses on how experts define the goals of the bachelor curriculum and how they evaluate AI's impact on those goals. Key considerations include academic integrity, skill development and the longer-term role of "manual" coding competence in a field where tools continuously automate parts of the workflow.

By exploring these perspectives, this thesis aims to map how experts connected to the Computing Science bachelor at Radboud University interpret the impact of AI tools on education in their field. The focus is on their views and reasoning and how they believe AI changes what students should learn, how they understand threats and opportunities for learning and assessment and where they personally position themselves on the spectrum between integration and restriction along with a literary review following the Rapid Review Method. The next chapter introduces related work to ground these perspectives. These consist of how AI may affect conceptual understanding and teaching practices, the rise of "vibe coding" as a new way of producing code and earlier major technological shifts that offer historical parallels for the current debate. Building on this foundation, the research chapter outlines the interview approach used to collect perspectives and the rapid review method, followed by a results chapter that organizes the main themes and points of agreement or disagreement. The thesis closes with a conclusion that summarizes the findings in relation to the research question and reflects on their implications for the bachelor program.

Chapter 2

Related Work

2.1 Educational evolution

Education has repeatedly had to renegotiate what counts as a 'basic skill' whenever a new tool made parts of schoolwork faster, easier or more automated. A classic example is the widespread availability of calculators. Meta-analytic research in mathematics education suggests that calculator availability does not automatically harm learning outcomes and may coincide with improved achievement and attitudes, particularly when instruction and assessment are designed with calculators in mind rather than treating them as a loophole[13].

A later wave came with the internet and specialized 'answer engines'. Tools such as WolframAlpha made it possible to obtain correct answers, sometimes including worked solutions, quickly. This again questioned whether something is a legitimate learning aid or does it bypass intended practice, which could cause a lack of learning due to the ease of usage[8]? This pattern is now visible in policy and review literature on educational technology, which stresses that learning effects depend on how tools are integrated and which teaching and assessment practices surround them[26].

These earlier transitions present a possibility for a pattern. Is generative AI mainly the next step in a familiar sequence or does it change educational conditions so much that existing approaches to assessment and skill development require deeper reconsideration?

2.2 Vibe coding

Within computing education specifically, generative AI introduces a new concept as students can now produce working code from natural-language prompts, even when they cannot yet explain how the code works or why it is correct. A term that has gained popularity for this style of development is 'vibe coding'. This describes an AI-dependent workflow where code is

produced through iterative prompting and quick trial-and-error rather than through careful manual construction. Vibe coding is not only a more accessible way for the average person to code, but could also make coding significantly faster for those who are already trained[15].

Within this form of education many courses aim to develop more than just the ability to produce a running program. Students are typically expected to learn to debug systematically, know about correctness and complexity and they need to be able to justify design choices. If AI systems can generate solutions quickly, some traditional programming assignments could shift from assessing implementation skill to assessing tool use or the ability to recognize plausible-looking outputs and or prompting skills.

Vibe coding may have a place in Computing Science with courses that aim to help students understand AI Assisted development. This could include the upside of fast paced coding and the downside of relinquishing some control over your code quality.

2.3 Impact of AI on understanding and learning goals

A central concern in the AI-in-education debate is the difference between *performance* (producing a 'correct' answer) and *understanding* (being able to explain, generalize, and apply knowledge independently). As mentioned in the above paragraph students are expected to learn many things beyond simple code generation. Generative AI can potentially support some of these aforementioned goals by providing explanations or alternative solutions that students can compare.

At the same time, a key risk discussed in recent literature is the overreliance of students and cognitive offloading. AI can perform the planning and reasoning steps where learners may skip the mental work that builds actual competence. This mainly happens when tasks are completed by accepting outputs rather than using AI as a scaffold[27]. In computing education specifically these concerns extend beyond plagiarism. They include shallow learning, reduced debugging practice and skill gaps that may only become visible later in the curriculum if learning objectives are not drastically changed[2].

2.4 Generative AI in education: effects on learning and study practices

The surge of large language model (LLM) tools such as ChatGPT has started many empirical research on how generative AI influences learning in higher education. Recent meta-analysis generally report positive average effects on student outcomes when ChatGPT is used as part of an instructional intervention, while also emphasizing that effects vary strongly by context and implementation. For example, a meta-analysis of 51 experimental studies published between November 2022 and February 2025 reports a large positive effect on learning performance (Hedges' $g = 0.867$), moderately positive effects on learning perception ($g = 0.456$) and higher-order thinking ($g = 0.457$). These were made with moderators including course type, learning model and intervention duration[28]. Systematic review and meta-analysis focusing specifically on experimental evidence (69 studies published between 2022 and 2024) concludes that ChatGPT interventions overall improve academic performance and motivational outcomes. These are associated with higher-order thinking behaviour, while also reducing mental effort. These papers highlight that, while ChatGPT may have a positive influence on learning goals, education should be adjusted[6].

Research also highlights that generative AI changes how students approach academic work. A systematic review of studies on ChatGPT in higher education (20 studies between January 2023 and March 2025) shows recurring themes such as perceived usefulness and efficiency. These come alongside concerns about overreliance, ethical ambiguity, content accuracy and academic integrity[9]. Internationally generative AI is seen as both an opportunity and a risk, emphasizing the need for AI literacy and responsible use practices rather than treating AI solely as a prohibited shortcut[26].

Related work suggests a consistent pattern in education. Generative AI can potentially support learning under certain instructional conditions, but it also introduces risks that are closely tied to how tasks are designed, how use is guided, and whether assessments can still properly measure required skills. These examples are based on many different studies, which still allows for uncertainty regarding its application in education of Computing/Computer Sciences.

2.5 Generative AI specifically in IT/CS education: effects on programming, assessment, and skill development

The rise of generative AI is impactful in IT and Computing Science because many core student objectives can be generated directly from prompts. Research on this problem has therefore focused not only on whether outcomes improve, but also on what changes in learning processes in this form of education.

A systematic literature review of generative AI published in late 2025 shows a rapidly expanding and diverse research landscape in computing science education. This review was made with studies considering multiple 'computing' topics. It reports both opportunities and challenges, which is similar to other education as mentioned above[1]. Similarly, another review of programming education studies (40 empirical studies) emphasizes that 'successful integration' tends to depend on teaching strategies, assessments and integration processes. It again flags over-reliance as a recurring concern, particularly regarding foundational programming logic and higher-order thinking[20]. In a similar direction, a systematic review specifically centered on harms and unintended consequences in computing education had concerns that again go beyond academic integrity alone. It foresees risks such as distorted feedback loops, shallow understanding masked by plausible outputs and skill gaps that may only become visible later in the curriculum[3]. These possible harms/consequences go further than challenges for other educations, which does create more risk for the computing science studies. However, this does not mean it should automatically be banned from these studies, as it might still be actively used.

Overall, current related work portrays generative AI in IT/CS education as a domain where the main challenge is not overreliance, but the alignment between learning goals and what assessments actually measure when code generation is becoming easier. This current situation leaves a rather large question to be answered, as is the question of this research. Education is different everywhere and those connected to the Computing Science bachelor at Radboud University have a better understanding of what could work and what would not. Change starts when questions are asked and those willing to listen take action.

Chapter 3

Research

This research uses two complementary methods. First, a pseudo-anonymous interview study was conducted with individuals currently or previously involved in teaching, advising or otherwise influencing the Bachelor of Computing Science at Radboud University. These interviews were used to explore how relevant actors currently view the role of generative AI in education and more specifically, within the Computing Science bachelor. Second, a rapid review was conducted to provide a structured overview of recent literature on generative AI in higher education and computing science education. Together, these methods allow the thesis to combine in-depth qualitative perspectives from the local context with broader insights from recent academic literature.

3.1 Interview study

A qualitative interview study was conducted to explore how individuals involved in the Bachelor of Computing Science at Radboud University currently view the use of generative AI in education. Because the research aimed to understand participants' perspectives, experiences and considerations in depth, a semi-structured interview design was chosen[24]. This format allowed the interviewer to prepare a consistent structure across interviews while still leaving room for participants to elaborate freely and introduce their own interpretations and concerns.

3.1.1 Participants

Ten potential participants, who currently or previously worked at Radboud University in teaching, advisory or otherwise influential roles within the Bachelor of Computing Science, were approached for inclusion, of which five agreed to participate. The final sample consisted of four male participants and one female participant. To remain unbiased the participants were not

informed about which other individuals participated in the study. To protect their privacy, the participants were pseudonymised. The participants were contacted via e-mail or asked in person.

3.1.2 Procedure

The interviews were held in private rooms, in the participants' office or a reserved room on the campus of Radboud University, as to retain the pseudonymity of the participants. The interviews were held only once per participant and lasted between approximately 75 and 105 minutes. Participants were offered the choice of conducting the interview in English or Dutch of which four chose to converse in Dutch and one in English. The interviews were recorded, after consent, for the purpose of transcription.

The interviews were designed such that the main questions were phrased as neutrally as possible. The interviews were held in a conversational manner, allowing substantial room for open interpretation and elaboration, as long as participants addressed the topic of the relevant interview section. This structure was chosen to encourage participants to explain not only their stance on generative AI, but also the reasoning and considerations underlying that stance.

The interview guide consisted of three main sections. These focused on participants' views and knowledge regarding generative AI in daily life and society more broadly, in higher education and specifically within the Bachelor of Computing Science at Radboud University. This sequence was designed to first establish broader perceptions of generative AI before gradually narrowing the discussion towards the educational and local context.

Although the interviews allowed for open conversation, guiding questions were used to ensure that each interview covered comparable themes.

Generative AI in daily life and society

This section explored participants' general stance and knowledge regarding the role of generative AI in everyday life and society more broadly. The guiding questions were:

- Would you want to see its role in society develop differently?
- Do you think generative AI is here to stay and likely to increase in use or do you expect its role to diminish over time?

Generative AI in education and higher education

This section focused on participants' views regarding the use of generative AI in education, and particularly in higher education. The guiding questions were:

- What is your stance on the use of generative AI by students?
- What is your stance on the use of generative AI by professors?
- Do you think generative AI can be used properly in education?
- Would your view change if Radboud University were to provide a private institutional AI environment, similar to examples at other Dutch universities such as VU Amsterdam?

Generative AI in the Bachelor of Computing Science at Radboud University

The final section narrowed the conversation to the context of the Computing Science bachelor. The guiding questions were:

- Do you use generative AI yourself in relation to your teaching or work?
- Are students in your courses allowed to use it, or do you follow general Radboud policy?
- Do you think generative AI can be meaningfully combined with your course or curriculum, or does it undermine too much of what students are supposed to learn?
- If so, do you think your course should be adapted?
- Would you potentially change your own course or the broader bachelor curriculum in response to generative AI?

Each interview ended with the question whether participants believed generative AI should, in the short term and in the long term, be banned from or integrated into the Bachelor of Computing Science at Radboud University.

3.1.3 Ethical considerations

All participants were informed about the purpose of the study, that participation was voluntary and the intended use of the interview data within this thesis. Participants provided informed consent before the interviews were recorded. To protect participant privacy, all interviews were pseudonymised during transcription and identifying information was removed from the dataset. Participants were not informed about which other individuals participated in the study. Audio recordings and transcripts were stored and generated locally on password-protected devices and were only accessible to the researcher to ensure data would remain local. One participant requested to review their transcript after transcription and no participants requested removal of statements or withdrawal from the study. As four interviews were

conducted in Dutch and one in English selected quotations used in this thesis were translated into English by the researcher. Care was taken to preserve the original meaning and phrasing of the participants as accurately as possible. The study was conducted in accordance with the ethical guidelines for qualitative research at Radboud University. No sensitive personal data beyond professional perspectives and experiences related to generative AI in education were collected.

3.1.4 Analysis

The interviews were analyzed using deductive thematic analysis[5]. Deductive thematic analysis is a qualitative analysis method in which coding and interpretation are guided by predefined themes (*a priori*) derived from existing theory, research questions, or prior findings.

The broader themes used in this study were informed by the research question, interview guide and findings from the rapid review, which resulted in five predefined themes:

- Perceptions of generative AI
- Generative AI in higher education
- Generative AI in the Bachelor of Computing Science
- Curriculum adaptation
- Banning versus integrating generative AI

The interviews were first transcribed using aTrain, which is a local and therefore offline system. The transcripts were then cut to exclude personal details and imported into ATLAS.ti, which was used to manage and code the data. Coding was conducted by one data coder. Codes were then compared to the predefined themes, while allowing for nuance and variation within each theme.

Following coding, coded fragments were revisited and interpreted within the broader themes in order to identify recurring patterns or notable differences across participants. Quotes were included to strengthen these findings. Participants did not provide feedback on the findings.

3.2 Rapid Review

To complement the qualitative interview study, this thesis includes a rapid review of the recent literature on generative AI in higher education and computing science education. A rapid review is a form of review that accelerates a traditional systematic review by streamlining or omitting selected steps, while maintaining a systematic and transparent approach. This method

is particularly suitable when time and resources are limited and a timely overview of available evidence is needed[12].

The rapid review serves three purposes in this thesis. It aims to map the main reported benefits, risks, and implementation considerations of generative AI in higher education, to inform the interview focus by identifying recurring themes in empirical research, and to support interpretation of the interview findings through combining them with published evidence.

3.2.1 Review question and objectives

The rapid review is guided by the following review question:

"What does recent literature report about the educational impact and governance of generative AI in higher education, particularly in computing science education?"

To apply this question, the rapid review pursues four objectives:

1. Summarize impacts of generative AI on learning processes and educational outcomes.
2. Identify challenges for assessment practices, including academic integrity and validity.
3. Describe implications for fairness, and student support needs.
4. Summarize the solutions and recommendations in the literature for using or regulating generative AI in higher education.

3.2.2 Eligibility criteria

Eligibility criteria were defined prior to searching to ensure a transparent and reproducible selection process. Studies were included if they met the criteria regarding publication characteristics, context, phenomenon of interest, and focus of outcomes.

Publication date and language

Studies published from 1 January 2022 onward were eligible for inclusion, as they coincided with the rise of public LLM usage. Publications written in English or Dutch were included, as these are the languages used in this thesis and the researched bachelor of Computing Science. Studies in other languages were excluded due to feasibility constraints.

Context and population

The review focused on higher education contexts. Studies situated in computing science, programming, or software engineering education were considered a priority subgroup, but studies from other higher education domains were also eligible if they addressed the review question.

Problem of interest

Studies were eligible if they examined the use or implications of generative AI and/or large language model tools in relation to teaching, learning, assessment creation or completion, or the governance of generative AI in higher education. Studies on educational technology not involving generative AI/LLMs were excluded.

Focus of outcomes

The focus of the rapid review is to provide an overview of recent neutral/empirical findings regarding the use of generative AI in higher education. Findings can reflect positive, negative, or mixed effects, depending on what was reported in the included studies.

Study design and publication type

Empirical studies using quantitative, qualitative, or mixed-methods designs were included. In addition, systematic reviews, scoping reviews, and meta-analyses focusing on generative AI/LLMs in higher education were eligible. Conceptual papers, editorials, and opinion pieces without a described method or empirical data were excluded. Peer-reviewed journal articles and peer-reviewed conference papers were included. Studies without accessible full text were excluded.

3.2.3 Information sources

The literature search was designed to be feasible within the time frame of this thesis while remaining transparent and reproducible. The search was conducted on Google Scholar as it contains a widely available and large set of finished literature often meeting the specified publication type.

3.2.4 Databases and search strategy

Searches were conducted in Google Scholar in April 2026. The publication period was restricted to studies published from 2022 onward. Language restrictions (English or Dutch) were applied during screening, as the filtering for this was not properly consistent.

3.2.5 Search queries

To keep the search process consistent and manageable, three predefined searches were performed. Two searches focused on generative AI in higher education in general (with different emphases), and one search focused specifically on computing science/programming education.

1. **Higher education (general):** *"generative AI" + "higher education"*
2. **Higher education (assessment/integrity focus):** *"generative AI" + "higher education" + "assessment"*
3. **Computing science / programming education:** *"generative AI" + "programming education"*

3.2.6 Procedure

After the search strategy and eligibility criteria had been established, the rapid review was conducted. Across the three search strings, 90 records were identified. After removal of duplicates, 73 records remained for screening. These records were assessed on the basis of their titles, abstracts, and bibliographic details to determine whether they met the predefined eligibility criteria. Following this initial screening stage, 25 studies were selected for full-text review. After full-text assessment, 11 studies were included in the final sample.

The 14 studies excluded after full-text review were omitted for several reasons. Some were too narrowly focused on a single application of generative AI in education, limiting their relevance to the broader review question. Others combined computing science education with engineering education in ways that made it difficult to isolate findings relevant to the present study. Additional studies were excluded because their evidential value was limited, for example due to very small sample sizes, or because their primary focus was on policy implementation rather than on the educational impact or governance questions central to this thesis.

As the review was conducted by a single researcher, no inter-rater agreement could be calculated.

3.2.7 Analysis

After study selection, the 11 included papers were analyzed using a narrative synthesis approach[21]. As the rapid review was based on three predefined search queries, the studies were first grouped according to the query through which they were selected. This resulted in three analytical groups. Four studies originated from the "generative AI" + "higher education query",

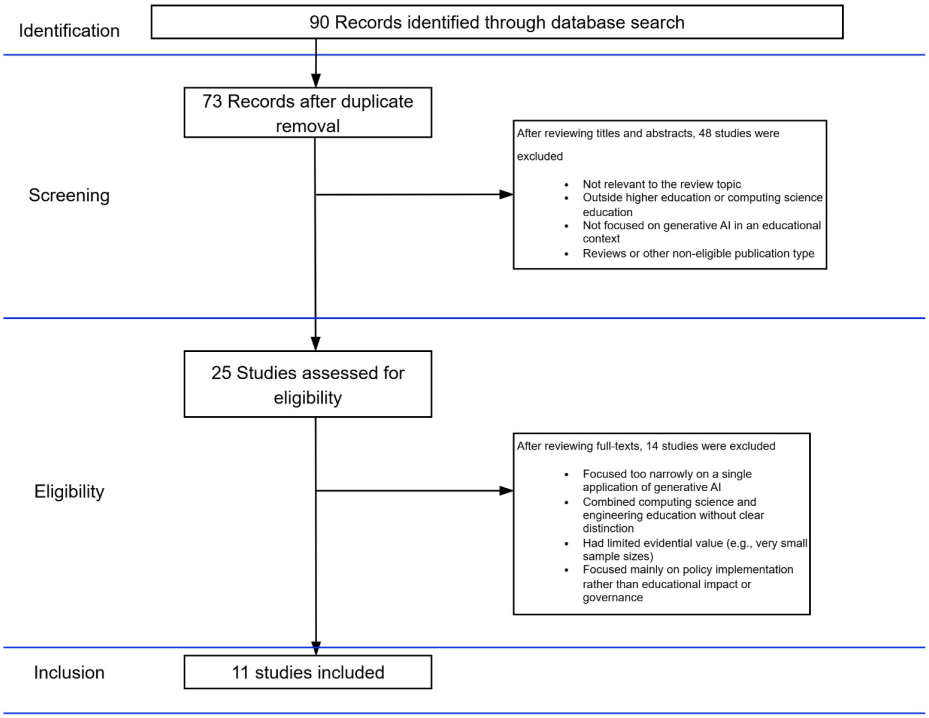

Figure 3.1: Rapid Review procedure.

three studies from the "generative AI" + "higher education" + "assessment" query, and four studies from the "generative AI" + "programming education" query.

Within each group, the studies were read and compared on their reported findings and conclusions. Particular attention was paid to recurring themes, such as reported effects on learning and the suggested role of generative AI in education. Differences in emphasis or conclusion between studies were also noted where relevant.

The findings were synthesized in two stages. First, each query group was summarized separately by identifying the main recurring conclusions across the included studies. Second, the grouped findings were compared across all three categories in order to identify broader patterns in the literature as a whole. This approach made it possible to move from query-specific findings to an overall interpretation of how recent literature reports on the educational impact and governance of generative AI in higher education, particularly in computing science education.

3.3 Trustworthiness

To strengthen the trustworthiness of this study, the criteria proposed by Lincoln & Guba were considered. These include credibility, transferability, dependability and confirmability[19].

3.3.1 Credibility

Credibility is considered the confidence of the "truth" of the findings. Credibility was primarily supported by methodological triangulation. The interview findings were considered alongside the rapid review findings, allowing the localized perspectives of the interviews to be compared with the broader patterns identified in recent academic literature.

This combination strengthened the interpretation of the findings, as themes identified in the interviews could be related to existing research on generative AI in higher education and computing science education.

3.3.2 Transferability

Transferability is the extent to which findings may be applicable in other contexts. The interview findings are inherently context-specific, as the study focused on individuals involved in the Bachelor of Computing Science at Radboud University. The findings should therefore not be interpreted as broadly representative of higher education institutions as a whole.

However, detailed descriptions of the institutional context, participant roles, interview structure and analytical process were provided to allow readers to assess the relevance of the findings to comparable contexts. The rapid

review included studies from multiple higher education settings, increasing the relevance of the reviewed literature.

3.3.3 Dependability

Dependability concerns the consistency and transparency of the research process. Dependability was supported by clearly documenting the research procedures, interview structure, rapid review procedure and the analytical approach used throughout the study.

Predefined themes and guiding questions helped maintain a relatively consistent structure across interviews, while still allowing flexibility for participants to elaborate on their own perspectives. The coding and analysis process was systematically conducted with predefined themes and documented using ATLAS.ti.

Within the rapid review, dependability was supported through the use of predefined eligibility criteria, predefined search queries and a transparent study selection procedure. However, as both the rapid review and interview analysis were conducted by a single researcher, no inter-rater agreement or formal inquiry audit could be established.

3.3.4 Confirmability

Confirmability regards the idea of neutrality in the research and how the study is shaped by the participants and not researcher bias, motivation or interest. As the interviews, rapid review, coding and interpretation were conducted by a single researcher, complete neutrality cannot be assumed.

Confirmability was supported through reflexivity by acknowledging the interpretative role of the researcher and the potential influence of the researcher's background within the context on the research process. A partial audit trail was maintained, to further increase confirmability, through documentation of the search strategy, eligibility criteria, the thematic structure and analytical decisions throughout the study.

Finally, direct quotations were included throughout the findings section to illustrate the connection between participants' responses and the interpretations made by the researcher.

Chapter 4

Results

4.1 Interviews

4.1.1 Findings

The interview analysis was informed by five predefined thematic areas derived from the research question and a preceding rapid review, which are general perceptions of generative AI, generative AI in higher education, generative AI in the Bachelor of Computing Science, curriculum adaptation, and the question of banning versus integrating generative AI. Within each theme codes were developed connecting direct quotations to the relevant themes.

Perceptions of generative AI

Participants broadly described generative AI as an unavoidable and increasingly influential development, while simultaneously expressing concerns regarding its reliability and impact. AI was frequently framed as a potentially valuable tool, but always considered to be very risky.

All participants emphasized that generative AI is unlikely to disappear, describing it as a technology that is "becoming more and more important" and a "keeper". While all assumed generative AI to stay, they were quite skeptic. One participant characterized the current state of AI as "a mess", pointing to issues such as hallucinated references and unreliable outputs. They addressed the concerns of generative AI taking over certain jobs or other tasks about which most, again, were skeptical. One person stated they do not believe generative AI would be used as a direct aid or assistant of sorts, as people want to remain in control: "I do not think people would want that, I don't think people would want to lose so much control."

A recurring concern was that users misunderstand the nature of generative AI. Participants noted that AI is often treated as a search engine, which they noted as expected since many models seem designed to look like it, despite functioning fundamentally different. As one participant ex-

plained, this creates a "false sense of certainty" where outputs are assumed to be factually correct even when they are not. Some accepted that generative AI could potentially generate better and faster product output, but were still extremely negative regarding the use of it: "but I wont deny that if you can produce a better product using AI that I understand your usage of AI in your area. However, would I hold it against you? Absolutely. You are helping the complete and utter destruction of our society."

These findings suggest that generative AI is not merely seen as a neutral tool, but as a potentially disruptive technology that challenges existing assumptions about online information and product creation. However, they acknowledged that it will remain relevant.

Generative AI in higher education

Participants highlighted both opportunities and risks associated with the use of generative AI by students. While AI was recognized as potentially useful for tasks such as brainstorming or refining text, concerns were raised about its impact on student learning processes. That refining text or letting generative AI create text hindered creativity and quality, "I think it is about how engaged your brain is with what you are currently trying to write." and writing it yourself forces you to mentally engage with it thus reconsidering everything that you wrote, "Writing is an essential part of your thinking process." One participant noted that students who directly input assignments into AI tools, destroy any learning opportunities as quoted: "aren't learning anything", while another argued that over-reliance on AI resembles a form of unproductive cheating, where students do not meaningfully engage with the content.

Participants acknowledged that students are likely to use AI regardless of institutional rules, particularly in situations involving time pressure or deadlines and that institutions should consider that possibility. For example stating that "We need to keep evaluating what we want students to learn, how we want them to learn and how realistic our wants are compared to reality." and "The learning process needs to be protected." This highlights the issue of not always being able to fight AI usage and therefore potentially having to adapt to secure quality of education.

These findings indicate that the challenge for higher education is not simply whether generative AI should be used, but how its use can be aligned with learning, especially focusing on ensuring that it does not lower the quality of education.

Generative AI in the Bachelor of Computing Science

In the context of the Bachelor of Computing Science generative AI was discussed primarily in relation to programming skills and the development of

computational thinking. Participants consistently emphasized that effective use of AI in programming requires prior knowledge stating that "to effectively use AI you sort of need to be a specialist in your field". Generative AI was therefore not framed as a replacement for programming education, but as a tool that may only become useful once students have developed sufficient understanding to evaluate its output.

One participant explained that in programming education there are clear differences between students who use generative AI critically and students who simply copy the assignment into an LLM. This can be seen by comparing students with programming expertise to those without as experienced students are better able to evaluate the output because they know, "when you are good at programming you know how to research it, how to critically research it" and can assess whether a program looks as expected or as they stated: "you try to test it and check if the program matches your expectations of what it should look like". This was strongly related to the idea of quality awareness. One participant argued that "recognizing quality becomes more and more important". Without this awareness, students who cannot yet program are, as one participant stated: "at the mercy of AI". This suggests that generative AI use in the Computing Science Bachelor does not remove the need for programming knowledge, but actually increases the importance of being able to judge code quality and to understand the fundamentals behind it.

This concern was also reflected in teaching practices where some courses restrict the use of generative AI to ensure that students develop core competencies independently. One participant explained that programming is: "a skill that you only acquire by messing around with it". Similarly, difficult programming concepts were described as things students sometimes need to "struggle with a bit", because those moments of struggle are part of the learning process.

At the same time, participants did not reject all forms of generative AI integration. There are considerations for a course *AI Assisted Development*, as an example of a more careful approach. One participant emphasized that this course is placed later in the curriculum because students must first develop "programming skills, logical thinking and algorithmic thinking". The same participant explicitly stated that the course "will not simply be vibecoding". This indicates that generative AI integration is being considered, but only after students have acquired the foundational skills needed to use it responsibly.

Overall, participants viewed generative AI in Computing Science as potentially useful, but only under specific conditions. Students first need to develop programming knowledge, critical thinking and quality awareness. In case they do not learn these generative AI risks becoming a substitute for learning rather than a support for learning. One participant summarized that AI may not always directly undermine learning in every context, but

in many foundational settings "it is so much easier to just be able to say, let's not do it."

Curriculum adaptation

Participants widely agreed that the rise of generative AI requires changes to the curriculum, particularly with regard to assessment practices. A central issue identified was the continued reliance on product-based assessment. Several participants noted that, "we have to work towards process based testing", as it is increasingly difficult to determine whether a submitted product was created by a student or generated by generative AI or as one participant stated "Our way of assessment is very product oriented. And that we cannot hold, because you cannot see if that product was made by AI or a student." This undermines the validity of traditional assessment methods, which focus primarily on these final outputs and not the process leading up to it. This approach of process based testing focuses how students arrive at their answers. Students have to include their usage of generative AI, their reasoning within the assignment and show their ability to evaluate outputs critically. Suggested approaches included tracking workflows, incorporating reflections, such as oral exams, and focusing on the development of critical thinking skills.

One participant did note that they would, if allowed, not adjust their courses stating that "I do not think certain courses will ever be changed, but I think our field will come to the conclusion: No, learning this form of programming is an essential skill." and that "in the courses that I teach it seems very clear that AI only undermines the learning goals".

These findings suggest that curriculum adaptation should be considered. However, it should not focus solely on restricting generative AI use, but on redefining what constitutes evidence of learning. At the same time, it should be accepted that some courses should be considered to remain the way they are currently and to more strictly handle generative AI.

Banning versus integrating generative AI

The question of whether generative AI should be banned or integrated into education was the final question and dilemma across the interviews. Rather than reaching a clear consensus, all participants described this issue as a complex problem, that should be dealt with in the relatively short term. One participant stated "This is not a black and white answer, this is a gray area".

Banning generative AI, particularly in the early stages of education, was seen as a way to protect the development of foundational skills. On the other hand, participants argued that integration is necessary, as generative AI is already widely used in professional contexts and by students. One

participant explicitly framed this as a dilemma between completely prohibiting generative AI usage and teaching students "to use it responsibly", also again referencing the potential of the *AI Assisted Development* course. This reflects a broader uncertainty within higher education regarding how to balance safeguarding learning processes and preparing students for future professional environments. Only one participant was incredibly strict regarding integration and would preferably consider it fraud stating: "It should just be treated as academical fraud and that the course assessment is of such a level that we can simply filter out 99% of those that do use it".

Both long term and short term were considered for answering this question, but the answers remained relatively the same as most did not want to try and "predict the future" in the long term.

Taking all points into the account the findings indicate that neither full prohibition nor unrestricted use is considered sufficient. Most participants pointed towards the need for a balanced approach in which students first develop core competencies and are subsequently taught how to use generative AI in a critical and responsible manner.

General findings

The interviews show that participants generally did not view generative AI as something that is simply good or bad. It was described as a development that is becoming too important to ignore, while at the same time bringing major risks for education. Across the interviews concerns about overreliance and the possible undermining of student learning returned repeatedly.

A strong pattern was that participants were less concerned with generative AI existing at all and more with the way it may replace important parts of the learning process. Programming and struggling with difficult material were repeatedly described as essential parts of education. AI was mainly seen as problematic when it takes over these parts too early or too fully.

Most participants did not argue for a complete rejection of generative AI in education. They generally recognized that AI is already part of professional life and will likely remain so. For that reason, many argued that students should at some point be taught how to use it critically and responsibly. This was usually seen as something that should only happen **after** students have first developed the necessary core skills themselves.

The interviews point towards a careful and selective approach with a slightly negative edge. Full prohibition was not seen as a sufficient long-term answer, but complete integration was also rejected.

The interviews therefore suggest that the impact of LLM use on education is mainly seen as double-sided. Generative AI may have value in later stages of learning and in professional contexts. Its current use however, was often seen as risky in relation to the protection of meaningful learning processes. For the Radboud Computing Science curriculum, the consequence

of integration would therefore not merely be the addition of AI-related content, but a broader need to reconsider how learning goals and assessment practices are aligned.

4.2 Rapid Review

4.2.1 Findings

The findings of the rapid review are organized according to the three analytical groups derived from the predefined search queries. This structure makes it possible to first examine the literature separately within each predefined category, before drawing a broader conclusion. The first subsection discusses studies on generative AI in higher education, the second focuses on studies concerning generative AI and assessment in higher education, and the third addresses studies on generative AI in programming education. The final subsection draws these findings together into an overall narrative synthesis.

Generative AI in Higher Education

The four studies selected for this group[11, 10, 16, 17] present generative AI as a development with major implications for higher education, but not as a purely positive or negative. In the selected papers generative AI is associated with several potential benefits like personalised learning and more support for students and staff. Francis et al. describe generative AI as a tool that can deliver substantial effiency gains and enhanced educational outcomes, while also emphasising that these benefits depend on very careful implementation. Similarly, Farrelly and Baker point to the potential value of AI for international students as a personalised language tutor and a potential afforable non human real-time support. Lee et al. also show that some educators already use AI as a practical tool for instance to test assessment tasks.

The studies consistently stress substantial risks especially in relation to academic integrity and overreliance on AI systems. Academic integrity is the most prominent of the concerns, but the papers show that this concern is more complex than simple student cheating. Francis et al. discuss the difficulty of distinguishing student-generated and AI-generated work. Farrelly and Baker add that AI detection tools are not reliable enough to serve as sole evidence of misconduct and that false accusations may especially affect international students and students writing in English as an additional language. They also mention bias and unequal access in relation to paid subscriptions Francis et al. argues the digital divide noting that access to more powerful AI tools may depend on paid subscriptions and reliable internet access.

The four studies converge on the view that higher education cannot simply ignore the presence of generative AI. The papers argue for an active response to address the situation. This includes clearer policy development, stronger AI literacy among both staff and students and adjustments to teaching and assessment practices. The selected studies do not advocate either adoption or outright rejection, but rather careful and responsible integration. Francis et al. recommend transparent policies, revised assessment designs and GenAI literacy programmes for both staff and students. Kurtz et al. provide a more explicit set of integration strategies. which includes for example raising awareness of disruptive change, training faculty and changing assessment practices. Lee et al. support this need empirically showing that less than a quarter of surveyed staff felt adequately equipped by their university to deal with AI, while more than three quarters indicated that they wanted support.

Taken together, the four studies in this category suggest that generative AI is best understood as a disruptive but potentially valuable development in higher education. Its impact depends less on the technology itself than on the extent to which institutions are able to shape its usage through adaptation and clear governance. Albeit assuming proper and fair usage of students, as is with all forms of education.

Generative AI in Higher Education Assessment

The three studies selected for this group[31, 29, 25] show a fairly clear shared view that generative AI is putting pressure on traditional forms of assessment in higher education. More conventional assessments such as essays, short-answer tasks, and take-home assignments are described as becoming less reliable when students can use AI tools to answer these problems with relatively little effort. Xia et al. state that tools such as ChatGPT can support students in completing tasks such as essay writing, proposal writing, and take-home examinations, which raises concerns about cheating and academic integrity. Similarly, Weng et al. argue that traditional assessment methods do not operate effectively in the GenAI era, because there is a mismatch between GenAI-integrated learning practices and traditional forms of assessment such as written essays and exams. This creates obvious concerns around academic integrity, but also around whether these assessments still properly measure student understanding and learning.

These studies are not entirely negative regarding AI usages and assessment. They suggest that higher education needs to rethink what assessment is supposed to do and what it should actually capture. Xia et al. identify possibilities such as self-assessment, immediate feedback and more diverse assessment methods, while also emphasising the need for teachers to design tasks that focus on critical thinking and creativity. Weng et al. distinguish between traditional assessment and GenAI-incorporated assess-

ment. In their review examples of GenAI-incorporated assessment include asking students to critique ideas and solutions generated by ChatGPT or demonstrate communication and collaboration skills through group-based tasks. Smolansky et al. similarly tested adapted assessment prompts in which students had to analyse and improve AI-generated essays or evaluate AI-generated code rather than simply produce a final answer themselves. Across the three papers it is clear that assessment should move towards more direct type of assessment. This includes approaches such as oral assessment, checking AI generated output and tasks that focus more on reasoning, application, and reflection than on producing a final written product alone.

Another recurring point in these studies is that the issue is again not only about cheating. The papers also connect assessment to a broader shift in learning outcomes. Skills such as critical thinking, creativity, self-regulated learning, AI literacy, and career-related competencies are repeatedly mentioned as becoming more important in a GenAI-influenced learning environment. Assessment is therefore not only described as a way to measure learning, but also as something that shapes what kinds of learning higher education should now prioritize. Weng et al. explicitly identify two new or refocused learning outcomes in GenAI-facilitated learning environments, which are career-driven competencies and lifelong learning skills. These include skills such as critical thinking and problem solving. Smolansky et al. also show that students and educators recognise the need for assessment reform although students were more hesitant than educators about adapted assessments because of concerns such as a possible loss of creativity.

The three studies in this category suggest that generative AI is pushing higher education assessment away from more traditional output-based formats and towards more adaptive, authentic, and skill-oriented approaches. The overall conclusion is that assessment reform is necessary, not only to reduce misuse of AI, but also to better match higher education with the kinds of knowledge and competencies students will need in an AI-influenced context. This grows only more relevant with more and more integration of AI into daily life and professional practice.

Generative AI in Programming Education

The four studies selected for this group[7, 4, 32, 23] show that generative AI is already starting to change programming education in clear ways. Across these papers GenAI is mainly presented as a tool that can support students through many different ways, for example code generation or feedback, during the programming process. Deriba et al. shows that this is already happening in their sytematic review where this was concluded from over 37 studies. Prather et al. describe opportunities such as using LLMs as pro-

gramming assistants in various ways and generators of teaching materials. This suggests that AI tools could be used to reduce some of the barriers that often make programming difficult, especially for less experienced learners.

The studies also connect GenAI use to motivation and engagement. Yilmaz et al. found in an experimental study with undergraduate programming students, that those who used ChatGPT during programming practice scored significantly higher than the control group on computational thinking skills, programming self-efficacy, and motivation. Boguslawski et al. similarly found that students were already making extensive use of LLMs and that these tools could improve learner autonomy and sense of competence by offering more personalised learning support, while also reducing negative emotions such as frustration or insecurity. GenAI can in this sense function as a low-threshold support tool that allows students to utilize a sort of personal teacher without the fear of negative judgement.

However, the studies are not completely positive about this development. A recurring concern is that students may become too reliant on AI-generated answers and therefore skip important parts of the learning process themselves. Prather et al. warn that students may use AI-generated code that they do not understand or be come over reliant on LLM's. Deriba et al. also identify concerns about over-reliance on GenAI tools and students accidentally accepting faulty answers. Boguslawski et al. add that although LLMs can reduce frustration and support motivation, they may also encourage students to accept generated content uncritically or avoid the struggle that is sometimes necessary for learning programming properly. This makes generative AI not only a learning support tool, but also a challenge for programming assessment and skill development.

Another point that comes back across these studies, is that programming education may need to shift its focus somewhat. Skills such as understanding code and evaluating AI-generated output, while also knowing how to use these tools critically are described as becoming more important. Prather et al. argue that the traditional approach of asking students to write many small programs may become less reliable because LLMs can solve many introductory programming tasks with simple prompts. They therefore suggest that computing education may need to place more emphasis on problem-solving, code or program understanding and the responsible use of AI tools. Boguslawski et al. also distinguish between resisting AI use and embracing it by teaching students to collaborate with AI, focus on code reading and critique and use the tools in ways that prepare them for future professional practice.

The papers still make clear that human support, structure, and guidance remain necessary as AI cannot fully replace the role of educators in supporting learning and motivation, as of yet. Boguslawski et al. explicitly argue that current LLMs cannot adequately provide or replace social

support, which remains a key factor in learner motivation and Prather et al. similarly report that students and instructors agreed that GenAI tools cannot replace human instructors.

Taken together, the four studies in this category suggest that generative AI in programming education should be seen as both useful and disruptive. It can support student learning, motivation and confidence, but it also creates clear risks if students start depending on it too much or use it without real understanding. The overall conclusion is therefore that these tools should not simply be rejected, but they should be used carefully in ways that still protect core programming skills, critical and logical thinking and meaningful learning.

General Findings

The studies included in this rapid review show that current literature is not necessarily positive or negative. The literature is broadly consistent in suggesting that generative AI is becoming too important to ignore in higher education, while acknowledging its benefits and downsides. Across the three groups of studies AI is repeatedly described as offering clear opportunities for learning and educational support, while at the same time, these same studies consistently point to serious risks. These risks are centered around academic integrity or unequal access, improper usage of its functionality and the potential weakening of critical/logical thinking skills.

The general pattern across the literature is therefore one of cautious acceptance. The reviewed studies do not support a complete rejection of AI in education, but neither do they support uncritical adoption. They suggest that AI should be integrated carefully and deliberately, with attention to clear policies, assessment design and the protection of meaningful learning outcomes. The broad conclusion across the reviewed papers is that the central question is no longer whether AI belongs in higher (programming) education, but how it can be integrated in ways to increase AI literacy while also preserving educational quality.

Chapter 5

Discussion

This study found that generative AI is viewed neither as a tool that should be fully banned nor as a tool that should be broadly integrated into the Radboud Computing Science curriculum. Instead, both the interview findings and the rapid review point towards a careful and selective approach. This is in line with existing literature, which similarly presents generative AI as both promising and disruptive in higher education[11, 10, 16, 17]. Previous studies have shown that generative AI can support learning and student motivation, but also creates various risks related to overreliance, assessment and reduced meaningful learning[28, 6, 31, 29]. The present findings mainly support this broader pattern, while adding that these risks become especially relevant in Computing Science education because AI can directly produce code and explanations that resemble these intended learning products of the curriculum[7, 4, 23].

5.1 Strengths and limitations

A defining strength of this study is that it combines two complementary sources of evidence. The rapid review provides an overview of recent literature on generative AI in higher education and computing science education, while the interviews add more detailed insight into how these broader developments are understood within the specific context of the Radboud Computing Science curriculum. This combination made it possible to compare local perspectives with wider patterns in the literature.

At the same time, several limitations should be considered. The interview study was based on a small sample of five participants from one specific educational context. The findings should therefore be carefully interpreted as it will likely not be a perfect representation of all Computing Science staff at Radboud University or of higher education more broadly. Both the interview coding and the rapid review were conducted by a single researcher, meaning that no inter-rater agreement could be established and that bias,

even if attempted to reduce as much as possible, will always be there. The rapid review was also limited by its use of predefined search queries and one main search database, which means that relevant studies may have been missed. Finally, because generative AI is developing quickly, the findings of this study should be understood as a reflection of a rapidly changing situation rather than as a final conclusion about the role of generative AI in education for all time.

5.2 Recommendations

Based on these findings, several recommendations can be made for higher education focusing on Computing Science. First, foundational programming courses should remain careful with the use of generative AI, especially when the use of generative AI would replace the practice needed to develop these programming skills through hardship. Second, generative AI should be integrated selectively and preferably later in the curriculum, when students have developed enough knowledge to critically evaluate AI-generated output. The Radboud University has already taken steps in this aspect with the current development of the AI Assisted Development course. Third, assessment should move beyond final products alone and include more attention to process, reasoning and justification, for example, if realistically possible, through oral explanation or process-based assessment. These recommendations do not imply that generative AI should be fully rejected, but that its use should be explicitly connected to learning goals and the protection of core skills, which is in line with the findings of the research.

Chapter 6

Conclusions

This research aimed to answer the question: *"What is the impact of usage of LLMs on education and AI governance, and what are the consequences for integration of AI usage in the Radboud Computing Science curriculum?"* This was done by combining a deductive thematic analysis of semi-structured interviews with participants connected to the Bachelor of Computing Science and a rapid review of recent literature on generative AI in education.

The findings show that LLM usage already has a substantial impact on education and AI governance. Both the literature and the interviews point to possible benefits, such as learning support and efficiency, but also to serious risks. These include overreliance, unreliable output and less meaningful learning processes. This means that action is needed in the relatively short term.

The findings do not support either full prohibition or broad integration of AI into the Radboud Computing Science curriculum. Instead, they point towards a selective approach. AI should not be ignored, but it should only be integrated where this can be justified in relation to the learning goals related to the subject. This suggests that generative AI partly follows an existing pattern of education adapting to new tools, but that its ability to directly produce code and explanations makes its impact on Computing Science education more fundamental and therefore does in fact require deeper reconsideration.

In conclusion, the impact of LLMs on education is significant, but their integration into the Computing Science curriculum should remain carefully limited. The challenge for higher education is therefore not simply whether to allow AI, but how to respond to it without undermining the core purposes of education, which also applies to the Bachelor of Computing Science.

Bibliography

[1] Friday Joseph Agbo et al. Computing education using generative artificial intelligence tools: A systematic literature review. *Computers and Education: Artificial Intelligence*, 2025.

[2] Russell Beale. Computer science education in the age of generative ai, 07 2025.

[3] S. Bernstein, A. Rahman, N. Sharifi, A. Terbish, and S. MacNeil. Beyond the benefits: A systematic review of the harms and unintended consequences of generative ai in computing education. In *Proceedings of ACM (SIGCSE-related venue)*, 2025.

[4] Samuel Boguslawski, Rowan Deer, and Mark G. Dawson. Programming education and learner motivation in the age of generative ai: student and educator perspectives. *Information and Learning Sciences*, 126(1-2):91–109, 07 2024.

[5] B Crabtree and W Miller. A template approach to text analysis: Developing and using codebooks. in b. crabtree. 1999.

[6] Ruiqi Deng, Maoli Jiang, Xinlu Yu, Yuyan Lu, and Shasha Liu. Does chatgpt enhance student learning? a systematic review and meta-analysis of experimental studies. *Computers & Education*, 2024.

[7] Fitsum Deriba, Ismaila Sanusi, Oladele Campbell, and Solomon Oyelere. Computer programming education in the age of generative ai: Insights from empirical research. *SSRN Electronic Journal*, 01 2024.

[8] Vincent Dimiceli, Andrew Lang, and Leighanne Locke. Teaching calculus with wolfram—alpha. *International Journal of Mathematical Education in Science and Technology*, 41:1061–1071, 12 2010.

[9] Izzet Dos. A systematic review of research on chatgpt in higher education. *The European Educational Researcher*, 2025.

[10] Tom Farrelly and Nick Baker. Generative artificial intelligence: Implications and considerations for higher education practice. *Education Sciences*, 13(11), 2023.

[11] Nigel J. Francis, Sue Jones, and David P. Smith. Generative ai in higher education: Balancing innovation and integrity. *British Journal of Biomedical Science*, Volume 81 - 2024, 2025.

[12] Chantelle Garritty, Gerald Gartlehner, Barbara Nussbaumer-Streit, Valerie J. King, Candyce Hamel, Chris Kamel, Lisa Affengruber, and Adrienne Stevens. Cochrane rapid reviews methods group offers evidence-informed guidance to conduct rapid reviews. *Journal of Clinical Epidemiology*, 130:13–22, 2021.

[13] Ray Hembree and Donald J. Dessart. Effects of hand-held calculators in precollege mathematics education: A meta-analysis. *Journal for Research in Mathematics Education*, 17(2):83–99, 1986.

[14] Andrej Karpathy. https://x.com/karpathy/status/1886192184808149383, february 2025.

[15] Andrej Karpathy. 2025 LLM year in review, 2025. Blog post; includes the term "vibe coding".

[16] Gila Kurtz, Meital Amzalag, Nava Shaked, Yanay Zaguri, Dan Kohen-Vacs, Eran Gal, Gideon Zailer, and Eran Barak-Medina. Strategies for integrating generative ai into higher education: Navigating challenges and leveraging opportunities. *Education Sciences*, 14(5), 2024.

[17] Daniel Lee, Matthew Arnold, Amit Srivastava, Katrina Plastow, Peter Strelan, Florian Ploeckl, Dimitra Lekkas, and Edward Palmer. The impact of generative ai on higher education learning and teaching: A study of educators' perspectives. *Computers and Education: Artificial Intelligence*, 6:100221, 2024.

[18] Ze Li, Nowshin Arony, Ahmed Awon, Daniela Damian, and Bowen Xu. Ai tool use and adoption in software development by individuals and organizations: A grounded theory study, 06 2024.

[19] EG. Lincoln, YS. Guba. *Naturalistic Inquiry*. Sage Publications, 1985.

[20] Jemimah Nathaniel, Solomon Sunday Oyelere, Jarkko Suhonen, and Matti Tedre. Literature review on the integration of generative ai in programming education. *International Journal of Artificial Intelligence in Education*, 35:2724–2755, 2025.

[21] Jennie Popay, Helen Roberts, Amanda Sowden, Mark Petticrew, Lisa Arai, Mark Rodgers, Nicky Britten, Katrina Roen, and Steven Duffy. *Guidance on the conduct of narrative synthesis in systematic reviews: A product from the ESRC Methods Programme*. 01 2006.

[22] Dimitri Prandner. What do students use ai tools for? assessing students' use of ai tools in three typical study related scenarios, 2025.

[23] James Prather, Paul Denny, Juho Leinonen, Brett A. Becker, Ibrahim Albluwi, Michelle Craig, Hieke Keuning, Natalie Kiesler, Tobias Kohn, Andrew Luxton-Reilly, Stephen MacNeil, Andrew Petersen, Raymond Pettit, Brent N. Reeves, and Jaromir Savelka. The robots are here: Navigating the generative ai revolution in computing education. In *Proceedings of the 2023 Working Group Reports on Innovation and Technology in Computer Science Education*, ITiCSE-WGR '23, page 108–159, New York, NY, USA, 2023. Association for Computing Machinery.

[24] Ruslin, Saepudin Mashuri, Muhammad Sarib, Firdiansyah Alhabsyi, Hijrah Syam, and Ruslin Ruslin. Semi-structured interview: A methodological reflection on the development of a qualitative research instrument in educational studies. 02 2022.

[25] Adele Smolansky, Andrew Cram, Corina Raduescu, Sandris Zeivots, Elaine Huber, and René Kizilcec. Educator and student perspectives on the impact of generative ai on assessments in higher education. pages 378–382, 07 2023.

[26] UNESCO. Guidance for generative ai in education and research, 2023. Policy guidance document.

[27] Aniella Mihaela Vieriu and Gabriel Petrea. The impact of artificial intelligence (ai) on students' academic development. *Education Sciences*, 15(3), 2025.

[28] Jin Wang and Wenxiang Fan. The effect of chatgpt on students' learning performance, learning perception, and higher-order thinking: insights from a meta-analysis. *Humanities and Social Sciences Communications*, 12:Article 621, 2025.

[29] XIA Q. Gu M. Rajaram K. Chiu T. K Weng, X. Assessment and learning outcomes for generative ai in higher education: A scoping review on current research status and trends. *Australasian Journal of Educational Technology, 40(6), 37-55*, 2024.

[30] Wilson Kia Onn Wong. The sudden disruptive rise of generative artificial intelligence? an evaluation of their impact on higher education and the global workplace. *Journal of Open Innovation: Technology, Market, and Complexity*, 2024.

[31] Qi Xia, Xiaojing Weng, Fan Ouyang, Tzung Jin Lin, and Thomas K.F. Chiu. A scoping review on how generative artificial intelligence transforms assessment in higher education. *International Journal of Educational Technology in Higher Education*, 21(1):40, 2024.

[32] Ramazan Yilmaz and Fatma Gizem Karaoglan Yilmaz. The effect of generative artificial intelligence (ai)-based tool use on students' computational thinking skills, programming self-efficacy and motivation. *Computers and Education: Artificial Intelligence*, 4:100147, 2023.