



ANCORing Generative AI within the computing curriculum

Mark Zarb, Robert Gordon University

Martin Goodfellow, University of Strathclyde

ABSTRACT

Generative AI and Large Language Models have become ubiquitous across education and within higher education institutions. Emerging challenges include potential over-reliance on Generative AI, risks to academic integrity, and inequitable access: there is an urgent need for students to develop ethical, self-regulated and grounded learning practices in its use. This paper presents insights distilled from a survey of 14 computer science educators in the UK, and identifies the overarching importance of teaching responsibility and ethical implications of the use of AI to students. The ANCOR framework is presented as a method for teaching responsible Generative AI use, integrating ethical reasoning, real-world examples, and curriculum-wide approaches. It offers a novel contribution by providing both actionable teaching techniques and a conceptual approach for embedding ethical and responsible use of Generative AI tools across computing curricula, including guidance on ethics integration, contextualising relevant policies, developing ethical decision-making skills, addressing anthropomorphism, and using illustrative real-world cases.

Keywords: ethics, computing curricula, generative AI, large language models, instructor insights

Introduction

Generative AI (GenAI) and Large Language Models (LLMs) have become ubiquitous (Ciampa et al., 2023; Zyda, 2024). Once challenging-to-use and inaccurate, GenAI tools have grown more knowledgeable and dependable for practically every user level. These systems are now in the hands of millions of consumers around the world (Caetano, 2025). Advancements to deep learning algorithms increase learning impact and boost the power of GenAI and LLMs. Millions of data points have been used to train this technology, and so long as it is readily available, its use seems inevitable. Therefore, ethical policies and recommendations need to be developed in compliance with contemporary standards (Zarb et al., 2024). This paper presents techniques, derived from a survey of experienced instructors, for introducing the concept of the ethical use of LLMs and GenAI tools into the context of computing education.

The growing popularity of these tools has compelled many educational institutions to suggest guidelines for their use. The surge in the use of LLMs and GenAI by students and educators has introduced new challenges and opportunities for UK HE institutions, many of which have introduced new guidance and/or updated their policies and procedures on the use of such tools by incorporating relevant guidelines on assessment redesign (QAA, 2023b), maintaining academic quality (QAA, 2023b), and academic integrity (QAA, 2023a). Additionally, extensive work has been conducted on the impact of LLMs and GenAI in the teaching, assessment and learning of specific Computer Science topics, such as programming (Denny et al., 2024; Finnie-Ansley et al., 2022; Goodfellow et al., 2025).

As the use of LLMs and GenAI becomes commonplace, the question of ethical use and responsibility arises naturally. Existing work (Prather et al., 2023; Yan et al., 2024) has considered the ethics of students using GenAI to support learning, as well as educators using GenAI to support teaching, such as the generation of learning resources and examples. However, previous work hasn't considered how to embed the teaching of ethics and responsibility across the curriculum.

The remainder of this paper is structured as follows: the next section provides background on computing education, including educator context and curriculum design considerations. This is followed by a description of data collection and analysis, with a discussion on key findings. The paper then presents the ANCOR framework and its application for teaching responsible GenAI use, concluding with a discussion on implications and directions for future work.

Background

There has recently been a large shift in how artificial intelligence technologies (AI) have been perceived and used in society (Zyda, 2024). AI - particularly LLMs and GenAI - has evolved significantly, transitioning from a challenging, error-prone discipline to a ubiquitous tool across diverse fields and contexts, and is accessible to millions of users around the world at all levels of digital literacy (Ciampa et al., 2023), making it a crucial part of today's digital landscape. This has a perceived impact on all aspects of modern society, including education.

The advancements in deep-learning algorithms have vastly improved LLM performance, and there is now a wide-spread perception that they are 'reliable' (Beluzzi et al., 2024). LLMs are trained on millions of data points, and can now be used to offer all manners of services, including practical, effective solutions across a wide range of tasks, from content creation to problem-solving. As these technologies continue to evolve, ethical considerations surrounding their use become increasingly important, especially in educational contexts. Advanced AI assistants, defined as artificial agents with natural language interfaces, present both opportunities and ethical challenges. These assistants can plan and execute actions across various domains, aligning with user expectations. However, their deployment raises concerns about value alignment, safety, and potential misuse. In educational settings, it is crucial to address issues such as manipulation, trust, and privacy to ensure these technologies are used responsibly and equitably (Gabriel et al., 2024).

Impact on teaching, learning and curriculum

Manipulation, trust, and privacy are all particularly pressing issues in fields like education, and therefore the introduction of these tools has led to a necessary but sudden shift in educational paradigms, including assessment practices, curriculum design and teaching strategies. Several UK higher education institutions and policy makers have updated their policies to address the challenges posed by GenAI tools (Potter et al., 2023), incorporating guidelines on assessment redesign, maintaining academic quality, and ensuring academic integrity (QAA, 2023a, 2023b). These changes reflect a broader effort to integrate ethical use of AI into educational frameworks, where these tools can have profound impacts on both teaching and learning (Denny et al., 2024; Finnie-Ansley et al., 2022; Gabriel et al., 2024). Furthermore, concerns about privacy and data usage continue to be significant barriers to the widespread adoption of GenAI in educational settings. As a result, many institutions are revisiting their policies to ensure that the use of these tools aligns with ethical standards and educational goals (Smith et al., 2024; Zarb et al., 2024) - although these policies are increasingly changing to deal with a rapidly shifting landscape.

The use of these tools lead to a number of administrative concerns around the authenticity of student work, leading to increasing cases in false authorship and plagiarism. Students' dependence (or in some cases, the term coined has been 'addiction') on AI tools - particularly at the lower levels where confidence has not yet been built (Salah et al., 2024; Zhang & Xu, 2024) - could undermine their ability to learn and solve problems independently, especially if they unlearn the ability to understand underlying concepts and start to rely heavily on blindly trusting the solutions provided by AI (Becker, et al., 2023b) without being necessarily cognisant – or wholly understanding – about the implication and possibility of existing bias within these LLMs. Moreover, issues of access and equity become relevant and significant, as not all students have the same level of access to these tools. At the time of writing, one of the more popular GenAI tools has publicised a number of tiers, including a free tier with limited access, a \$20pcm tier advertised as “Plus” and a \$200pcm tier advertised as “Pro”, giving “the best answers to the highest questions”. This will inevitably create disparities in learning opportunities (Smith et al., 2024).

Within computing education, this surge in the use of GenAI tools and technologies has created new opportunities and challenges for students and educators, particularly regarding the notions of responsible and ethical use. Existing literature has explored the role of AI in supporting student learning, with particular attention to its potential to enhance understanding of complex topics, such as programming (Becker, et al., 2023b). These tools have been shown to be able to benefit students with varying levels of prior programming experience; for example by generating syntactically and logically correct code from natural language prompts, explaining error messages, debugging broken code, or generating code prompts for further work (Becker, et al., 2023a; Becker, et al., 2023b). This does, of course, raise concerns. It has been seen that students (and perhaps even educators) can become overly reliant on AI tools prior to contextualising or scaffolding foundational skills, which could then inhibit the development of critical independent problem-solving skills (Prather et al., 2023; Yan et al., 2024).

GenAI may have an impact on the entire computing curriculum – from shaping how we approach teaching fundamental programming to pedagogical approaches used within this discipline (Becker, et al., 2023b). Studies suggest that including GenAI may lead to significant changes in students' emotional intelligence (Keshishi & Hack, 2023), engagement, graduate attributes, and skills development (Mahon et al., 2024). These shifts demand immediate action, as educators and institutions must navigate the challenge of harnessing the opportunities presented by these tools while upholding high standards of academic integrity and ensuring meaningful learning outcomes.

Several studies have investigated the implications of GenAI tools on self-regulated learning in educational contexts. Research has shown that students' skills, such as self-efficacy, intrinsic motivation and effort regulation, can influence their adoption and effective use of AI tools for learning tasks (Mirriahi et al., 2025). Other work highlights how learner characteristics and the affordances of GenAI tools interact to enhance self-regulated learning capacities, emphasising the importance of user engagement (Wu & Chiu, 2025; Xu et al., 2025). While these tools can assist students in generating correct code or offering solutions, there is concern that they may interfere with the development of essential self-regulation strategies, and therefore impact on long-term learning. These strategies are critical for students to plan, monitor, and reflect on their problem-solving processes, all of which are vital skills for success in both academic and professional settings (Amoozadeh et al., 2024). It is imperative that the educational community guides the integration of these tools into curricula in ways that enhance learning without compromising academic integrity (Becker, et al., 2023a; Smith et al., 2024). There is a need for educational frameworks that not only incorporate AI tools into the learning process but also focus on teaching students about the ethical, societal, and legal implications of AI technologies. Without this, students may fail to grasp the potential

risks and responsibilities associated with these tools (Prather et al., 2023), especially as they learn when it is appropriate to use (and not use) these tools, and within what context.

While the integration of GenAI and LLMs into computing education offers significant benefits, including increased efficiency, enhanced learning experiences, and the development of new skills, it also raises substantial challenges. These challenges include concerns about student dependence on AI, issues of academic integrity, equity in access, and the ethical implications of AI tools. The growing body of literature on this topic underscores the need for comprehensive guidelines and policies that address both the opportunities and risks posed by these technologies. As the use of AI tools becomes more prevalent, it is crucial that educational institutions continue to explore ways to integrate these tools into curricula in a responsible and ethical manner, preparing students for the future of work while ensuring that core educational principles are upheld.

Data collection

Method

The authors ran a workshop at the ACM Computing Education Practice (CEP) conference in 2024, on the topic: *Large Language Models (LLMs) and Computer Science Education: How do we adapt our classes and practices?*

This study used an explanatory sequential mixed-methods approach (Creswell & Plano Clark, 2018), for providing both breadth and depth of insights from educators. A pre-workshop survey consisting of eight quantitative questions was distributed to all attendees, specifically targeted towards educators within UK higher education institutions, and aimed to explore how educators used GenAI within their practice. Due to CEP being a computer science education research conference, this excluded educators not involved in computer science education research and academics without teaching responsibilities. The survey questions were designed to capture educators' backgrounds, teaching context, current practices with GenAI, where and how its use is allowed in student work, and perceptions of pedagogical and ethical priorities, ensuring that each response informed the study's aim of understanding GenAI integration in higher education computing courses. The survey results were used to structure discussion points during the workshop, which provided qualitative insight, and responses from both are used as the basis of this position paper.

Ethical Clearance

This study makes use of data from human participants, and approval has been granted for the use of this data via the University of Strathclyde, which covers the use of this data, as well as its analysis, for research purposes. Furthermore, participants gave informed consent for the duration of the study and for the use of their data as outlined above, including its analysis for research purposes.

Participants

Fourteen instructors from six institutions across the UK and Ireland responded to the survey. All instructors had experience in teaching computing programmes. Their teaching experience was captured in Table 1,

with the majority of participants having 10+ years experience, closely followed by 3-5 and 5-10 years experience. Only one participant had less than three years' teaching experience.

Table 1: Participants' Teaching Experience

Teaching Experience	Percentage
0 – 2 years	7.1%
3 – 5 years	28.6%
5 – 10 years	28.6%
10+ years	35.7%

The levels taught by each participant can be seen in Table 2. The majority taught undergraduate classes, followed closely by both undergraduate and postgraduate classes. As this paper is written in a UK context, we refer to 'classes' or 'modules' as individual units of study within a larger degree programme, whereas the term 'course' refers to the degree programme as a whole.

Table 2: Participants' Teaching Level

Level	Percentage
Undergraduate	50%
Postgraduate	7.1%
Both	42.9%

The primary activities relating to their role can be seen in Table 3. As expected, all participants were involved in teaching. However, only 50% were involved in assessment design. The numbers decreased for module design, followed by course design. Those responsible for course design, as would be expected, were involved in all activities. The only other activity mentioned was research, by two participants (14.1%), but this is outside the scope of this paper.

Table 3: Participants' Primary Activities Relating to their Role

Primary Activities	Percentage
Teaching	100%
Assessment Design	50%
Module Design	42.9%
Course Design	35.7%
All of the above	35.7%

Results and initial discussion

The survey consisted of a number of Likert-scale questions in order to ascertain where and how the respondents were using GenAI within their curriculum. These responses will be captured in this section to set the scene.

Initially, respondents were asked how they use GenAI in their own practice. Five participants (35.7%) don't make use of any GenAI in their practice. This can be seen in Figure 1. There was no evident relationship between GenAI usage and experience.




Figure 1: Survey Respondents around Use of Generative AI

The results of what it is used for can be seen in Table 4. Unsurprisingly, the most common use was for assistance generating examples. Another comment mentioned using it for generating images. The least common use was for developing slides, which, from experience, is likely due to its tendency to create busy slides with “walls of text”. Interestingly, providing feedback wasn't common despite this being a major focus of computer science education research. However, this is more common for programming education (Cambaz & Zhang, 2024; Prather et al., 2023). Further studies would be required to determine if this was dependent on the types and topics of modules taught.

Table 4: Use of Generative AI

How do you use Generative AI in your practice?	Percentage
Generating Examples	57.1%
Developing Slides	7.1%
Developing Lecture Content	21.4%
Developing Assessments	35.7%
Providing Feedback	14.3%
Misc. admin tasks related to the role	28.6%
I do not use Generative AI in my practice	35.7%

Respondents were also asked where they allowed students to use Generative AI within the curriculum. These results are presented in Table 5.

Table 5: Where Generative AI is allowed for use in Student Work

Where do you allow student use of Generative AI in their work?	Percentage
Formative Assessments	57.1%
Summative Assessments	42.9%
Labs	57.1%
Tutorials	42.9%

ANCORing Generative AI within the computing curriculum

As a general rule, I do not allow students to use Generative AI in their work	21.4%
No restriction	7.1%
University policy is to allow use	7.1%
I have a guide outlining what is acceptable practice for different parts of my class	7.1%

It is interesting to note a disparity in the allowance of GenAI across different types of assessments. While a majority of instructors permit its use in formative assessments and labs, a number of respondents also indicate that they permit its application in summative assessments and tutorials. Interestingly, the majority of participants (5/7 or 71.4%) who are involved in assessment design also allow the use of GenAI, with only one of these participants mentioning university policy.

In addition, 3/5 or 60% of the participants that don't use GenAI in their practice allow their students to use it (as opposed to 2/9 or 22.2% who use it but don't allow their students to use it). However, 2/5 or 40% stated it was due to university policy, so this may not be the case otherwise. Further studies would be needed to determine if there was a relationship between educators who use GenAI in their practice and those that allow its use.

It is also worth noting the general low responses that point to existing institutional or class-level policies on the acceptable practice of these tools. This is confirmed in the follow up question with three participants mentioning university guidelines and one mentioning a module guide. This disparity raises questions about the role of institutional oversight and the autonomy of educators in regulating GenAI usage. This prompts considerations about the collaborative development of comprehensive guidelines that strike a balance between fostering innovation and upholding ethical standards in AI education.

When further prompted if guidance is provided for how students can use GenAI in their courses, only one person explicitly said no and the only other negative response has plans to change this. Others focused on general guidelines for its use, such as being aware of the negatives, trying not to become over reliant or even just trying to convince them not to use it. One required it to be used for assessment, whereas others stated that if it was used students had to be able to explain everything they produced and were only allowed to use constructs from the class or declare generated code via comments.

Table 6 showcases responses to the question: *How important do you think it is to teach students how generative AI works?*. Perhaps unsurprisingly, the data skews towards the top end of the scale.

Table 6: Perceived Importance of teaching how Generative AI works

How important do you think it is to teach students how generative AI works?	Percentage
1 (Not important)	0.0%
2	14.3%
3	21.4%
4	21.4%
5 (Very important)	42.9%

The lowest scores also don't use it in their practice but vary on allowing student use. However, the majority that don't allow student use still thought it was important to teach. Interestingly, 2/6 or 33% of the respondents who allow it to be used for summative assessment were neutral about teaching the students how it works. All others who use it for summative assessment ranked it as important.

An analysis of follow-up responses to this question indicates the integration of AI education should begin at a foundational level, being included as part of computer science concepts in general education, and further delved into in specialised technical courses at higher academic levels. While deep technical understanding of how AI functions may not be necessary for all students, comprehension of its implications, such as ethical use and biases in its implementation, is crucial.

Respondents argue that the perception of AI as a 'magic box' should be debunked, and students should be taught to approach it critically. This entails understanding its limitations and recognising it as a tool rather than a standalone solution. By doing so, students would be able to efficiently utilise AI as a support tool while remaining aware of its drawbacks and limitations.

This is seen to be especially important in computer science disciplines, where it is imperative for students to have a clear understanding of the inner workings of the technologies they are using. This knowledge would enable them to accurately assess the effectiveness and applicability of AI in various contexts, ensuring informed decision-making and responsible use.

Table 7 shows the Likert-scale breakdown of the responses to the question *How important do you think it is to teach students about the ethics of using generative AI?* The majority of respondents rated the importance of teaching students about the ethics of these tools as "most important", with none indicating that this was an unimportant topic. 42.9% thought ethics was more important than teaching students how it worked, with only one respondent (7.1%) thinking ethics was less important. The remaining 50% thought they were of equal importance.

Table 7: Perceived Importance of teaching the ethics of using Generative AI

How important do you think it is to teach students about the ethics of using Generative AI?	Percentage
1 (Not important)	0.0%
2	0.0%
3	28.6%
4	7.1%
5 (Very important)	64.3%

When elaborating on why, reasons were given such as these tools being unavoidable but they should be treated just like other software tools, a required skill for professional practice. When asked how they would teach the ethics most participants mentioned using examples and focussing on the negatives of GenAI, such as hallucinations and climate cost, as well as how one might use it unethically.

This question in particular prompted discussion at the follow-on workshop. The workshop was structured as a group discussion guided by the open-ended questions from the survey to encourage talking points. The subsequent discussion was captured by the authors, who noted the emergence of the topics of transparency, responsibility and ethical use as focal points. This discussion was distilled into a number of points which will form the basis of the following section.

ANCOR: Teaching responsible use of generative AI tools

The survey results have highlighted the pressing need for educators to emphasise the importance of using GenAI tools ethically and responsibly. While the survey provided valuable insights into the current practices

surrounding GenAI usage in computing education, it also underscored the imperative for educators to guide students in understanding the ethical implications of their use of these tools. This section explores strategies for integrating discussions on responsible GenAI use across the curriculum, drawing from both the survey findings, follow-on workshop discussions, and existing literature.

- **A - Beware Anthropomorphism:** Students have been observed to anthropomorphise GenAI tools, attributing human-like qualities that can lead to inaccurate mental models and undermine critical evaluation of AI. This connects directly to broader concerns in the literature about students becoming overly reliant on AI tools prior to contextualising or scaffolding foundational skills, which can inhibit critical problem-solving (Prather et al., 2023; Yan et al., 2024). A tendency to anthropomorphise may accelerate this dependence by encouraging students to ‘trust’ the tool as they would a human expert. To mitigate this, instructors should highlight the limitations of AI assistants, and engage students in discussions about the potential pitfalls of viewing AI as human-like beings. Students should be empowered to adopt a critical and sceptical mindset when interacting with these tools, allowing them to develop informed perspectives and make responsible decisions. This includes the understanding that AI is a tool that processes data to produce outcomes, and not an intelligent agent. The Raspberry Pi Foundation has some helpful hints on how to avoid anthropomorphism¹.
- **N - Integrate Discussions Across the Curriculum:** Several programmes and courses restrict discussions about responsibility, ethics and ethical use of tools to a single module that emphasises students’ professional responsibilities (Goetze, 2023; Grosz et al., 2019). However, a pressing need to talk about responsibility and ethical use in a range of classes and contexts across the curriculum has now emerged. For example, consideration of privacy could be factored into classes on database systems paired with discussion of privacy regulations such as GDPR. In the context of AI, discussion of climate impact could also be included when determining appropriate approaches and methods. This would help students to start to think at a critical level, not only about how to use these tools to suit their needs, but also on when might or might not be appropriate. This approach reflects wider calls for computing curricula to address the full range of impacts that GenAI may have on students’ engagement, graduate attributes and skills development (Mahon et al., 2024).
- **C - Contextualise Policy/Situation:** At the time of writing, many higher education institutions and professional bodies are issuing statements and policies on the appropriate and ethical use of GenAI (ACM Publications Board, 2023; Potter et al., 2023). It is important to highlight the broader sense of policy implications to students (QAA, 2023a, 2023b). Students should not only understand where to find policies and regulations that affect them and their work, but also be willing to engage in ongoing discussions and developments in the field. In addition, it is likely that these guidelines will need to be subject or even module-specific: the rules that are in place for a business module might not translate to a computer science context. Many institutions are revisiting policy to ensure alignment with ethical standards and educational goals (Smith et al., 2024; Zarb et al., 2024). As a practical first step, an adaptation of the Student Guide in Appendix D of (Prather et al., 2023) could be provided at the module or course level, offering students a concrete and contextualised starting point.

¹ <https://www.raspberrypi.org/blog/ai-education-anthropomorphism>

- **O - Obvious Real-World Examples:** Experiential learning draws heavily on the use of real-world examples. Within the context of GenAI, potential ethical issues and consequences can be framed as case studies, helping students appreciate the relevance and importance of ethical considerations when using these technologies, particularly in complex projects such as capstones. The literature highlights the value of such grounding: case studies addressing bias, privacy, and human judgement provide tangible contexts for students to apply ethical reasoning, while also illustrating inequities in access and the potential for misuse (Becker, et al., 2023b; Smith et al., 2024). By engaging with these real-world scenarios, students can better understand the societal and ethical implications of GenAI use, reinforcing responsible decision-making. Instructors should provide opportunities for students to discuss and debate their position, so that they can explore different perspectives on the topic. These examples could focus on areas such as bias, privacy and human judgement. For example, one of the most widely known instances of bias was evident in the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) software, which was used to determine the likelihood of a defendant becoming a reoffender. This software contained racial bias, which resulted in it being far more likely to determine black defendants were more at risk of reoffending than their white counterparts. There are also games which can help students see how bias can be introduced into a dataset, such as Survival of the Best Fit (www.survivalofthebestfit.com). Alternatively, as a focus on human judgement, the class could discuss the lawyers that were fined after submitting a legal filing containing fake citations.
- **R – Responsibility Across the Curriculum:** Instructors should provide opportunities for students to develop skills for ethical decision-making in the context of GenAI, by providing them with a safe space to learn how to evaluate ethical dilemmas, weigh different considerations and make informed decisions that align with the principles of their field. This connects to concerns in the literature that, without the appropriate scaffolding, GenAI use may interfere with the development of self-regulation strategies (Mirriahi et al., 2025; Wu & Chiu, 2025; Xu et al., 2025). Building on case studies, students can be presented with a hypothetical situation or ongoing case for critical discussion; for example, they could be asked to consider the dangers of models being trained on sensitive data, taking into account privacy regulations and IP laws.

Conclusion

This paper presented a distilled set of discussion points from an audience of computer science educators, on how to better embed GenAI within the curriculum. The ANCOR framework addresses calls for educational approaches that go beyond teaching students to use AI tools, guiding them to understand the ethical, societal and legal implications of these technologies (Becker, et al., 2023a; Smith et al., 2024). The discussions focussed quite heavily on the importance of teaching responsibility and ethical use of these tools to students at various stages of their learning journeys, rather than treating it as an afterthought. As these technologies continue to rapidly evolve, educators must ensure that students are equipped with appropriate critical skills that allow them to adapt alongside these changes, both during their studies and once they join workplaces where these technologies are ubiquitous. These findings also have implications for universities and academic practice, highlighting ways to integrate ethical and responsible GenAI use throughout computing curricula and inform curriculum design.

Whilst the authors have started the process of embedding these techniques within their own curricula across two Scottish institutions, the impact of this is still being evaluated, and will be reported in a future paper, highlighting the need for further waves of investigation.

Biographies

Mark Zarb is an Associate Professor based within the School of Computing, Engineering and Technology at the Robert Gordon University. He is involved in both teaching and research, weaving both towards the development of courses and modules within the computing curriculum. His research work often focuses on improving educational practices in computing, with an emphasis on developing effective teaching methodologies and tools to enhance learning experiences for students in this field. He is a Senior Fellow of AdvanceHE and a Chartered IT Professional.

Martin Goodfellow is a Senior Teaching Fellow in Computer and Information Sciences at the University of Strathclyde, where he is Director of Undergraduate Teaching and Deputy Group Lead of the Computer Science Education Research Group. He is a Fellow of AdvanceHE and has over 10 years' experience of teaching computer science.

References

- ACM Publications Board. (2023, April). ACM policy on Authorship. Association for Computing Machinery. <https://www.acm.org/publications/policies/new-acm-policy-on-authorship>
- Amoozadeh, M., Daniels, D., Nam, D., Kumar, A., Chen, S., Hilton, M., Srinivasa Ragavan, S., & Alipour, M. A. (2024). Trust in Generative AI among Students: An exploratory study. *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1, SIGCSE 2024*, 67–73. <https://doi.org/10.1145/3626252.3630842>
- Becker, B. A., Craig, M., Denny, P., Keuning, H., Kiesler, N., Leinonen, J., Luxton-Reilly, A., Prather, J., & Quille, K. (2023a). Generative AI in introductory programming. Association for Computing Machinery.
- Becker, B. A., Denny, P., Finnie-Ansley, J., Luxton-Reilly, A., Prather, J., & Santos, E. A. (2023b). Programming Is Hard—Or at Least It Used to Be: Educational Opportunities and Challenges of AI Code Generation. *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1, SIGCSE 2023*, 500–506. <https://doi.org/10.1145/3545945.356975>
- Beluzzi, F., Condorelli, V., & Giuffrida, G. (2024). Does It Really Work? Perception of Reliability of ChatGPT in Daily Use. *Italian Sociological Review*, 14(10S), 625–655.
- Caetano, D. L. (2025). *AI and Sustainability: How the Use of GenAI Tools Influence Consumer Choice for Sustainable Products* (Master's thesis). Universidad NOVA de Lisboa, NIMS – Dissertações de Mestrado em Marketing Analítico.
- Cambaz, D., & Zhang, X. (2024). Use of AI-driven Code Generation Models in Teaching and Learning Programming: A Systematic Literature Review. *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1, SIGCSE 2024*, 172–178. <https://doi.org/10.1145/3626252.3630958>
- Ciampa, K., Wolfe, Z. M., & Bronstein, B. (2023). ChatGPT in education: Transforming digital literacy practices. *Journal of Adolescent & Adult Literacy*, 67(3), 186–195.
- Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (Third edition). Sage.
- Denny, P., Prather, J., Becker, B. A., Finnie-Ansley, J., Hellas, A., Leinonen, J., Luxton-Reilly, A., Reeves, B. N., Santos, E. A., & Sarsa, S. (2024). Computing Education in the Era of Generative AI. *Commun. ACM*, 67(2), 56–67. <https://doi.org/10.1145/3624720>
- Finnie-Ansley, J., Denny, P., Becker, B. A., Luxton-Reilly, A., & Prather, J. (2022). The Robots Are Coming: Exploring the Implications of OpenAI Codex on Introductory Programming. *Proceedings of the 24th Australasian Computing Education Conference, ACE '22*, 10–19. <https://doi.org/10.1145/3511861.3511863>

ANCORing Generative AI within the computing curriculum

- Gabriel, I., Manzini, A., Keeling, G., Hendricks, L. A., Rieser, V., Iqbal, H., Tomašev, N., Ktena, I., Kenton, Z., Rodriguez, M., & others. (2024). The ethics of advanced ai assistants. arXiv Preprint arXiv:2404.16244.
- Goetze, T. S. (2023). Integrating ethics into computer science education: Multi-, inter-, and transdisciplinary approaches. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1* (pp. 645-651).
- Goodfellow, M., Booth, R., Fagan, A., & Lambert, A. (2025). AutoMCQ - Automatically Generate Code Comprehension Questions using GenAI. In *Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 2* (pp. 737-738). Association for Computing Machinery. <https://doi.org/10.1145/3724389.3731266>
- Grosz, B. J., Grant, D. G., Vredenburg, K., Behrends, J., Hu, L., Simmons, A., & Waldo, J. (2019). Embedded Ethics: integrating ethics across CS education. *Communications of the ACM*, 62(8), 54-61.
- Keshishi, N., & Hack, D. S. (2023). Emotional Intelligence in the Digital Age: Harnessing AI for Students' Inner Development. *Journal of Perspectives in Applied Academic Practice*, 11(3), 172-175. <https://doi.org/10.56433/jpaap.v11i3.579>
- Mahon, J., Mac Namee, B., & Becker, B. A. (2024). Guidelines for the Evolving Role of Generative AI in Introductory Programming Based on Emerging Practice. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1* (pp. 10-16).
- Mirriahi, N., Marrone, R., Barthakur, A., Gabriel, F., Colton, J., Yeung, T. N., Arthur, P., & Kovanovic, V. (2025). The relationship between students' self-regulated learning skills and technology acceptance of GenAI. *Australasian Journal of Educational Technology*, 41(2), 16-33.
- Potter, J., Welsh, K., & Milne, L. (2023). Evaluating an institutional response to Generative Artificial Intelligence (GenAI): Applying Kotter's change model and sharing lessons learned for educational development. *Journal of Perspectives in Applied Academic Practice*, 11(3), 139-152. <https://doi.org/10.56433/jpaap.v11i3.582>
- Prather, J., Denny, P., Leinonen, J., Becker, B. A., Albluwi, I., Craig, M., Keuning, H., Kiesler, N., Kohn, T., Luxton-Reilly, A., MacNeil, S., Petersen, A., Pettit, R., Reeves, B. N., & Savelka, J. (2023). The Robots Are Here: Navigating the Generative AI Revolution in Computing Education. *Proceedings of the 2023 Working Group Reports on Innovation and Technology in Computer Science Education*, ITICSE-WGR '23, 108-159. <https://doi.org/10.1145/3623762.3633499>
- QAA. (2023a, May). Maintaining quality and standards in the ChatGPT era: QAA advice on the opportunities and challenges posed by Generative Artificial Intelligence. In QAA advice and resources on Generative AI. https://www.qaa.ac.uk/docs/qaa/members/maintaining-quality-and-standards-in-the-chatgpt-era.pdf?sfvrsn=2408aa81_10
- QAA. (2023b, July). Reconsidering assessment for the Chat GPT era: QAA advice on developing sustainable assessment strategies. In QAA advice and resources on Generative AI. https://www.qaa.ac.uk/docs/qaa/members/reconsidering-assessment-for-the-chat-gpt-era.pdf?sfvrsn=38d3af81_6
- Salah, M., Abdelfattah, F., Alhalbusi, H., & Al Mukhaini, M. (2024). *Me and My AI Bot: Exploring the 'Alholic' Phenomenon and University Students' Dependency on Generative AI Chatbots - Is This the New Academic Addiction?* (Version 3). Research Square. <https://doi.org/10.21203/rs.3.rs-3508563/v3>
- Smith, C. E., Shiekh, K., Cooreman, H., Rahman, S., Zhu, Y., Siam, M. K., Ivanitskiy, M., Ahmed, A. M., Hallinan, M., Grisak, A., & Fierro, G. (2024). Early Adoption of Generative Artificial Intelligence in Computing Education: Emergent Student Use Cases and Perspectives in 2023. *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1*, ITICSE 2024, 3-9. <https://doi.org/10.1145/3649217.3653575>
- Wu, X.-Y., & Chiu, T. K. (2025). Integrating learner characteristics and generative AI affordances to enhance self-regulated learning: A configurational analysis. *Journal of New Approaches in Educational Research*, 14(1), Article 10. <https://doi.org/10.1007/s44322-025-00028-x>
- Xu, X., Qiao, L., Cheng, N., Liu, H., & Zhao, W. (2025). Enhancing self-regulated learning and learning experience in generative AI environments: The critical role of metacognitive support. *British Journal of Educational Technology*, 56(5), 1842-1863. <https://doi.org/10.1111/bjet.13599>

ANCORing Generative AI within the computing curriculum

- Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., Li, X., Jin, Y., & Gašević, D. (2024). Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1), 90–112. <https://doi.org/10.1111/bjet.13370>
- Zarb, M., Brown, J. N. A., Goodfellow, M., Liaskos, K., & Young, T. (2024). Ethical Implications of Gen-AI and LLMs in Computing Education. *Proceedings of the 2024 on ACM Virtual Global Computing Education Conference V. 2, SIGCSE Virtual 2024*, 293–294. <https://doi.org/10.1145/3649409.3691074>
- Zhang, L., & Xu, J. (2024). The paradox of self-efficacy and technological dependence: Unraveling generative AI's impact on university students' task completion. *The Internet and Higher Education*, 65, Article 100978. <https://doi.org/10.1016/j.iheduc.2024.100978>.
- Zyda, M. (2024). Generative AI changes the world, maybe. *Computer*, 57(9), 118–123. <https://doi.ieeecomputersociety.org/10.1109/MC.2024.3416231>.