ENHANCING REFLECTIVE TEACHING IN HIGHER EDUCATION

GenAI in Higher Education

Redefining Teaching and Learning

Sam Illingworth and Rachel Forsyth

BLOOMSBURY

GenAI in Higher Education

Enhancing Reflective Teaching in Higher Education

Series Editors:
Paul Ashwin (Lancaster University, UK)
Margaret Blackie (Virginia Polytechnic Institute and State University, USA)

The *Enhancing Reflective Teaching in Higher Education* book series is committed to supporting the development of reflective teaching in higher education. Books in this series undertake in-depth explorations of contemporary issues related to the development of reflective educational practices and are informed by the latest research and thinking about teaching in higher education. They are designed around evidence-informed principles for effective educational practices and provide strategies to enhance day-to-day teaching in higher education. In doing so, they build on the foundations provided by *Reflective Teaching in Higher Education*.

Books in this series will explore ways of enhancing reflective teaching in areas such as Doctoral Education and Team-based Learning as well as through strategies such as the Scholarship of Learning and Teaching and Students as Change Makers.

Forthcoming in the series:
Reflective Practice on Doctoral Education, edited by Friona Hallett

GenAI in Higher Education

Redefining Teaching and Learning

Sam Illingworth and Rachel Forsyth

BLOOMSBURY ACADEMIC
LONDON • NEW YORK • OXFORD • NEW DELHI • SYDNEY

BLOOMSBURY ACADEMIC
Bloomsbury Publishing Plc, 50 Bedford Square, London, WC1B 3DP, UK
Bloomsbury Publishing Inc, 1359 Broadway, New York, NY 10018, USA
Bloomsbury Publishing Ireland, 29 Earlsfort Terrace, Dublin 2, D02 AY28, Ireland

BLOOMSBURY, BLOOMSBURY ACADEMIC and the Diana logo are
trademarks of Bloomsbury Publishing Plc

First published in Great Britain 2026

Series design by Grace Ridge
Cover image: stellalevi via Getty Images

A catalogue record for this book is available from the British Library.

A catalog record for this book is available from the Library of Congress.

ISBN: HB: 978-1-3505-3579-4
PB: 978-1-3505-3578-7
ePDF: 978-1-3505-3581-7
eBook: 978-1-3505-3580-0

Series: Enhancing Reflective Teaching in Higher Education

Typeset by Integra Software Services Pvt. Ltd.

For product safety related questions contact productsafety@bloomsbury.com.

To find out more about our authors and books visit www.bloomsbury.com and sign up for our newsletters.

Contents

Tables

Series Editors' Foreword

The *Enhancing Reflective Teaching in Higher Education* book series is committed to supporting the development of reflective teaching in higher education. Books in this series undertake in-depth explorations of contemporary issues related to the development of reflective educational practices and are informed by the latest research and thinking about teaching in higher education. They are designed around evidence-informed principles for effective educational practices and provide strategies to enhance day-to-day teaching in higher education. In doing so, they build on the foundations provided by *Reflective Teaching in Higher Education*.

Many higher education educators are unsure about how to respond to the possibilities and challenges of GenAI. *GenAI in Higher Education: Redefining Teaching and Learning* is designed to support reflective teachers in higher education whatever their current level of understanding of GenAI. It supports them to responsibly evaluate, integrate and govern GenAI so that it can enrich, rather than compromise, existing educational practices. Built around four foundational principles of Student-Centeredness, Trust, Relevance and Agency, it supports educators to use GenAI across the full teaching and learning lifecycle. This includes the design of inclusive activities, fostering student engagement, giving feedback and reflecting on educational practices. Each chapter explores practical ways to apply GenAI, helping reflective teachers in HE to make informed, creative choices that are aligned with their educational values. As a whole, the book will enable educators to integrate theoretical and practical applications of GenAI, anticipate and prepare for future GenAI needs and critically assess the ethical and practical implications of using GenAI in their educational practices.

Acknowledgements

Writing this book has been a genuinely collaborative process, and we are very grateful to the many people whose generosity and critical insight have shaped it along the way. We would especially like to thank the colleagues and students we have been learning alongside, whose curiosity, critique and lived experience continue to challenge and expand our thinking. And, of course, our thanks also go to the GenAI tools themselves – the sometimes surprising, often slippery companions we have been working with to better understand both the opportunities and the barriers they present to learning and teaching in higher education.

We are especially grateful to the editorial team at Bloomsbury, and to Mark Richardson in particular, for their encouragement, responsiveness and good humour throughout. Thanks also to series editors Mags Blackie and Paul Ashwin for their intellectual guidance and belief in the project from its earliest stages. Lund University library made the process of open access publishing seem very easy, and we hope that this availability gives more people the opportunity to reflect on and discuss the place of generative AI in higher education.

We would like to thank Kirsten Jack, Mirjam Glessmer and Tom Lowe for their detailed, generous and always kind feedback on the manuscript; your care and attention made a real difference. Our warmest thanks also to the anonymous reviewers whose thoughtful comments helped us refine the argument and sharpen the structure.

To everyone who has challenged, questioned and supported our thinking about GenAI and higher education – thank you. This book is richer because of you.

Sam and Rachel

1

Introduction

About This Book

The rapid emergence of Generative Artificial Intelligence (GenAI) presents higher education with unprecedented opportunities to augment and enhance learning. As educators, learners, technologists and institutional leaders, we find ourselves navigating uncharted territory. How can we evaluate, integrate and govern these rapidly evolving technologies responsibly, so that they enrich rather than displace the irreplaceable human elements of education? We have written this book to try and address these questions, to offer guidance for realizing GenAI's potential while upholding the values of thoughtful pedagogy and academic rigour.

This is not a book about technology but a book which aims to consider how these technologies can be used as part of four foundational principles that underscore every chapter, every example and every case study. Ultimately, we believe that higher education should be centred around:

1. **Student-Centeredness** – Ensuring that students can engage with valued forms of knowledge, quality, standards and expertise.
2. **Trust** – Fostering an open and trusting learning environment where students feel comfortable taking risks and expressing themselves authentically.
3. **Relevance** – Making learning relevant and meaningful by connecting taught concepts to students' lives and interests and to the wider social context.
4. **Agency –** Giving learners and educators autonomy and voice in shaping their higher education journeys.

Sam is a professor at Edinburgh Napier University in the UK, whose research and practice are focused on interdisciplinary studies and creative pedagogies. And Rachel, based in Sweden, is an expert in curriculum design, digital learning and assessment. Together, we blend our complementary experiences and perspectives to offer you, the reader, this guide to navigating the opportunities and challenges of integrating GenAI in higher education. In these pages, we aim to synthesize the latest evidence, case studies and emerging practices, combining theoretical grounding with practical frameworks you can apply in your own practice. Our goal is to empower you to evaluate GenAI tools thoughtfully, implement them effectively and govern them responsibly to enrich teaching and learning.

This book is designed for a wide readership with an interest in GenAI's potential influences on higher education, including:

- Educators exploring GenAI capabilities to enhance teaching practices, provide more personalized and adaptive learning, automate routine tasks and offer instant student feedback and support.
- Administrators and policymakers aiming to integrate GenAI tools via responsible protocols and governance that uphold educational values.
- Educational technologists considering specific GenAI tools and platforms to determine suitability for their institution's needs.

Whether you are an education leader implementing institutional strategy, an academic developer supporting staff adoption, a curious educator or a wary sceptic, this book offers multi-lensed perspectives into GenAI that empower thoughtful, strategic integration focused on human-centric educational values. In sharing our own experiences and those of others, we hope to collectively shape how GenAI tools can catalyse discovery, increase access and enrich learning without undermining the human relationships that must remain at the core of the higher education experience.

This book is designed to support you in using GenAI across the full teaching and learning lifecycle – from designing inclusive, GenAI-informed activities to fostering student engagement, giving feedback and reflecting on your own practice. Each chapter explores practical ways to apply GenAI, helping you make informed, creative choices that align with your own practice. By the end of this book, you will be able to:

1. **Integrate theoretical and practical GenAI applications –** build a solid understanding of GenAI's theoretical foundations alongside practical applications.

2. **Anticipate and prepare for future GenAI needs in higher education –** recognize emerging skills and knowledge required for GenAI's effective educational use.
3. **Critically assess ethical and practical implications of GenAI –** develop a thoughtful, critical approach to using GenAI responsibly and ethically.

We recognize that this is a rapidly evolving field, with new tools and capabilities emerging at pace. However, the pedagogical principles that underpin this book – student-centredness, trust, relevance and agency – are designed to remain valid regardless of which technologies come and go. Where we have used GenAI to help develop learning prompts or examples, this is clearly indicated in the text, so that readers can reflect on their role and value as part of an open and critical learning process.

Learning Outcomes

Chapter 2 introduces the concept of Generative AI (GenAI) and its potential implications for higher education. This chapter is designed to help you start thinking critically about GenAI's role in your current educational practices and how it might fit into your future goals. The learning outcomes focus on helping you establish a baseline understanding of GenAI and reflect on its potential applications in your context.

By the end of this chapter, you will be able to:

1. Brainstorm potential applications of GenAI in higher education, considering both opportunities and challenges in your specific context.
2. Set personal or institutional goals for how GenAI might support educational practices, student engagement or administrative processes.
3. Reflect on your current use (or non-use) of GenAI in teaching, learning, or management, and establish a baseline for future exploration.

These learning outcomes will serve as a framework for the exercises, allowing you to test your understanding of the material and to think critically about how GenAI might be integrated into your own teaching, learning or administrative contexts. By engaging with the exercises, you will be better prepared to consider GenAI's potential, while keeping a balanced perspective on its capabilities and limitations.

What Is GenAI?

Artificial intelligence (AI) refers broadly to systems that exhibit intelligent behaviours such as learning, problem-solving and perception. This encompasses many different types of AI systems and techniques like machine learning, deep learning, computer vision and natural language processing.

GenAI is a subset of AI focusing specifically on generating new content like text, images, audio and video. These systems are trained on large datasets to learn patterns and can then use that learning to create new, original outputs. Unlike systems that simply retrieve and display pre-existing information, GenAI can produce novel, human-like content. Two common types are text generators like GPT-4 that can write passages, stories or articles, and image generators like DALL-E that can create pictures based on text descriptions. The aim is to mimic human creativity and imagination. While the outputs are not always perfect, GenAI has improved greatly in recent years at tasks like writing, drawing, composing music and translating between languages. We provide a more detailed description of the technology in Chapter 2.

This opening chapter introduces the rapidly evolving landscape of GenAI tools in higher education. Subsequent chapters offer practical guidance grounded in theory and practice on crafting human-centric GenAI implementations tailored to local institutional contexts and challenges. We also provide an analysis of pivotal issues in the use of GenAI, including data ethics, privacy, algorithmic bias, the digital divide and responsible design principles.

We hope that this book empowers readers to harness emerging capabilities as amplifiers of human creativity, inclusion and discovery while ensuring educational values remain grounded and ever present.

The Evolution of GenAI

The twenty-first century has ushered in an era marked by the symbiotic relationship between technology and education. This period has been characterized by the rapid introduction of technologies in higher education, intended to enrich students' learning experiences and create flexible learning opportunities, building on students' supposed digital expertise. While chatbots, the early manifestations of AI, played their part, it is the comprehensive capabilities of GenAI that have the potential to reshape the contours of higher education. GenAI systems, transcending the limited functionalities of their

predecessors, are equipped to comprehend and respond to intricate human language nuances with a precision hitherto unseen.

The narrative of AI begins in the 1960s with the creation of ELIZA, a prototype natural language processing program developed at the Massachusetts Institute of Technology (MIT) in the United States by Joseph Weisenbaum (Sharma et al., 2017). The subsequent decades saw a proliferation of conversational agents, with the likes of PARRY, crafted by psychiatrist Kenneth Colby in 1972, and Jabberwacky, the brainchild of British programmer Rollo Carpenter, developed in 1988. While these systems marked significant progress in AI, they were mere waypoints on the journey to GenAI.

A pivotal juncture in AI evolution was the unveiling of A.L.I.C.E. (Artificial Linguistic Internet Computer Entity) in 1995. Developed by Richard Wallace, A.L.I.C.E. introduced the world to a pioneering approach that facilitated intricate conversational dynamics (AbuShawar & Atwell, 2015). This innovation garnered much acclaim, securing the Loebner Prize (an annual AI competition since 1991 where systems try to convincingly mimic human conversations) three times between 2000 and 2004 (Loebner, 2009).

As the internet's tendrils spread deeper into our daily lives in the 2000s, AI, and subsequently GenAI, began to redefine online interactions. The rise of social media platforms further amplified the relevance of AI, culminating in Facebook's launch of its chatbot in 2016, thereby catalysing a surge in AI development.

Such chatbots were able to parse simple human queries and suggest preformed simple answers based on fixed databases of material. For instance, chatbots could provide plausible responses to queries about ordering processes or stock levels in an online shop, and then trigger a request for a human intervention when the limits of its database had been reached.

A further notable development in the field of artificial intelligence is the advent of large language models (LLMs), which represent a significant leap forward in machine learning. LLMs, such as GPT-3 and GPT-4, have demonstrated an ability to generate human-like text by analysing and predicting language patterns from vast datasets. Their capacity to engage in more complex, contextually aware conversations offers new possibilities for both automated systems and personalized learning (Alsafari et al., 2024).

In November 2022, OpenAI released a publicly available GenAI model, ChatGPT. This software generated new text in response to user questions. It was quickly followed by the release of several other systems such as Google's BARD and then Gemini (Floridi, 2023), Microsoft's Bing Copilot, Anthropic's Claude and many more. These systems, backed by vast textual data, have the

capability to engage in human-like dialogues, answering intricate questions and fostering engaging conversations.

The Role of Technology in Higher Education

Higher education institutions have long been breeding grounds for innovation and progress, adapting to societal changes and technological advancements to better serve their students and educators. Over time, technology has played a crucial role in streamlining administrative processes, offering more efficient communication channels and expanding access to educational resources. For instance, Learning Management Systems (LMS) have simplified course management for instructors, while digital textbooks and online databases have made learning materials more readily available to students regardless of their geographical location. These advancements have both increased operational efficiency and fostered more flexible, student-centred learning environments, allowing for personalized pathways and self-paced study options (Bates & Sangra, 2011).

Now, GenAI is emerging as a tool which manufacturers promise could address some of the persistent challenges in higher education. As higher education continues tackling issues like rising costs, accessibility pressures and demand for diverse learning experiences, GenAI is said to offer new capabilities that could catalyse positive change if thoughtfully implemented, and so there is strong temptation to explore its potential. When deployed thoughtfully alongside educators' expertise and care, GenAI might become a transformative catalyst.

In 2023, the EDUCAUSE organization, representing over 2,100 educational institutions in the United States, conducted a 'Quick Poll' of its higher education members to understand how universities and colleges were adapting to the presence of GenAI tools and what kinds of applications they were exploring. The poll received 441 responses, which were categorized into four areas of use: Dreaming, Drudgery, Design and Development (McCormack, 2023).

- **Dreaming** encapsulates the creative and exploratory potential of GenAI. Respondents highlighted applications such as brainstorming new ideas, generating questions for academic inquiry and exploring innovative solutions to complex problems. This use of GenAI allows educators and administrators to experiment with possibilities that might not emerge through traditional methods, fostering a culture of innovation. For example,

educators could use GenAI to brainstorm interdisciplinary research ideas or refine course content to better engage students.

- **Drudgery** refers to the use of GenAI for automating repetitive, time-consuming tasks, freeing up staff to focus on more meaningful work. Examples include drafting routine administrative documents, managing email correspondence and summarizing lengthy reports. By addressing these tasks, GenAI not only reduces workload but also minimizes human error in routine processes. However, this area also raises concerns about dependency on automation and the potential for oversimplifying complex tasks, which may require nuanced human judgement.
- **Design** involves using GenAI to create materials such as presentations, visual aids and course resources. Respondents noted that this capability is particularly valuable for educators seeking to enhance the accessibility and engagement of their teaching materials. GenAI tools can quickly generate high-quality visualizations or tailor resources to diverse learning needs, enabling educators to focus on pedagogy rather than production. Nevertheless, ensuring that these designs meet institutional standards and inclusivity requirements remains a critical consideration.
- **Development** focuses on strategic and policy-oriented uses of GenAI, such as drafting institutional policies, developing strategic plans and preparing funding proposals. This category highlights GenAI's role in shaping the future direction of institutions. Respondents observed that while these tools can streamline complex processes, their effectiveness depends on effective oversight and governance to prevent misuse and maintain institutional integrity.

The Four 'D's framework offers higher education institutions a pathway to strategically experiment with GenAI technologies. Dreaming and Design present opportunities for creativity and innovation, while Drudgery and Development highlight practical applications that can address immediate needs. Understanding these distinctions can help institutions balance experimentation with risk management as they integrate GenAI into their operations.

McCormack suggests that institutional leaders use this framework to identify lower-risk areas for initial adoption while developing policies to govern higher-stakes applications. For example, automating routine administrative tasks (Drudgery) might require minimal oversight compared to drafting strategic policies (Development), which involves ethical and organizational complexities.

By categorizing GenAI use cases into these four areas, the EDUCAUSE Quick Poll provides a lens through which higher education can consider its future with

these emerging technologies – both as a tool for immediate efficiencies and as a driver of long-term transformation.

Building on these findings, GenAI's generative capabilities have the potential to contribute to various routine university processes, including content creation, customized learning resources and simulations. If developed responsibly with rigorous testing and oversight, such artificial outputs could significantly expand educators' toolkits. Moreover, the potential for analysing patterns in student data to create highly personalized learning experiences could benefit students by tailoring education to their strengths, needs and interests. This would be particularly valuable for marginalized student groups, often poorly served by one-size-fits-all educational models (Development, and sometimes Dreaming). Indeed, early evidence suggests that neurodiverse students are already beginning to use these tools to personalize their own learning (Malmström et al., 2023), and adaptive learning pathways may enhance engagement, satisfaction and outcomes for diverse learners.

These categories illustrate the wide range of potential applications for GenAI tools in universities and can serve as a useful framework for considering where to explore their use further. They may also help in prioritizing development resources toward areas where GenAI could offer the most value.

Since the publication of the EDUCAUSE Quick Poll, educational software developers have begun creating tools that extend beyond these initial categories, with a particular focus on enhancing students' experiences without increasing the workload for educators.

Some of the key promises being made by these developers include:

- Enhancing student support and engagement through personalized GenAI tutoring systems, virtual assistants and collaborative tools that make learning more interactive.
- Offering personalized learning experiences by analysing student patterns and adapting resources to match each individual's level, needs and interests, optimizing their development.
- Automating assessment and feedback to provide timely support, allowing educators to focus on other important activities.
- Assisting educators by streamlining administrative tasks, offering easy access to resources and handling frequent queries efficiently.

As with any commercial technology developed for broad, generic purposes, careful consideration is required to determine how GenAI should be adapted

within an educational context. Just as tools like Microsoft Word or PowerPoint have enhanced productivity without replacing the need for human oversight, GenAI cannot substitute for the human relationships, judgement and care that are essential to education and learning. Educators' roles in motivating, mentoring and designing thoughtful learning experiences tailored to students' diverse needs remain crucial. GenAI is not a cure-all for the challenges in higher education, such as the need to rethink curriculum design and assessment strategies. Rather, it should be viewed as a developing set of capabilities that, when combined with pedagogical expertise, might empower and enrich education for all.

Large-scale implementation of any educational technology carries risks, including concerns about data privacy, algorithmic bias, misinformation and over-automation. Successful implementation requires careful planning, thoughtful customization and innovative integration with existing practices, all while respecting the irreplaceable skills of educators. Ethical considerations and inclusion must also be prioritized in thinking about implementation. This has been true for earlier technologies, such as Microsoft Word and Learning Management Systems (LMS) or Virtual Learning Environments (VLE), and remains just as relevant for GenAI tools today.

The main difference with GenAI lies in its apparent power and the speed at which it has been adopted in society – factors that can sometimes distract from the central focus of education: the intention and impact of teaching, rather than the tools themselves. In later chapters, particularly Chapter 3, we will explore these issues in more detail. We will also explore GenAI techniques, practical applications, implementation strategies, ethical considerations and future directions in relation to higher education and to the four foundational principles we outlined earlier: student-centredness, trust, relevance and agency. In doing so we hope to provide insights and guidance for integrating GenAI's possibilities thoughtfully while heeding necessary ethical and pedagogical cautions.

How to Use This Book

This book provides a comprehensive yet accessible overview of GenAI technologies and their potential applications in higher education, paired with practical guidance on how to evaluate, select, implement and govern them effectively and ethically in your own practice. We offer the following as a

potential roadmap for navigating and making the most of this book with regard to your own learning and teaching journeys:

- Having read this introductory chapter, you should have some foundational knowledge regarding the evolution of GenAI systems and techniques. This context will illuminate the possibilities these rapidly advancing technologies introduce for transforming teaching, learning, assessment and student support.
- Proceed through the chapters systematically or else select topics of particular interest and relevance. The chapters are designed to stand alone, enabling you to focus on the most pertinent areas given your role and institutional context.
- Work through the reflective exercises and further reading at the end of each chapter. These activities encourage critical analysis and dialogue that will enrich your understanding, as well as giving opportunities to add to your teaching portfolio.
- Contribute your own experiences and insights to the growing community of higher education professionals navigating this technology shift by joining mailing lists and communities of practice. Sharing your successes and challenges with peers will accelerate collective learning.

Exercises

Each chapter in this book includes exercises to help you reflect on and apply the discussed concepts. You will often see prompts to use ChatGPT, as it is one of the GenAI tools we use most frequently, but feel free to substitute with other tools, such as Gemini, Claude, or similar alternatives. For Chapter 1, the exercises focus on brainstorming, goal setting and reflecting on current practices with GenAI, guiding you in critically considering how GenAI could shape your educational context.

Exercise 1.1: GenAI Brainstorming

Suggested time: 15 minutes.

Learning outcome: brainstorm potential applications of GenAI in higher education, considering both opportunities and existing challenges in your specific context. Try to think freely without censoring yourself regarding

your own instinctive preferences, ethics, practicalities and other aspects. We will come back and critically assess these ideas in later chapters.

Description: list three potential applications of GenAI that could align with your institution or programme's goals. Think about both teaching and administrative processes. Consider how these applications could address current challenges, enhance student learning or improve efficiency. Focus on practical cases that could be explored further.

Examples:

- Automating administrative tasks such as grading or attendance tracking.
- Developing personalized tutoring systems that adapt to individual students' needs.

Exercise 1.2: GenAI Goal Setting

Suggested time: 15 minutes.

Learning outcome: set personal or institutional goals for how GenAI might support educational practices, student engagement or administrative processes.

Description: reflect on your own goals as an educator, administrator or student, and identify three areas where GenAI could help you achieve these goals. Consider both immediate and long-term objectives, as well as potential challenges or barriers. How might GenAI assist in achieving these goals, and what limitations might you need to account for?

Examples:

- As an educator, I want to use GenAI to provide instant feedback to students, helping me offer more timely and constructive support.
- As an administrator, my goal is to reduce the time spent on repetitive queries, so I plan to integrate a GenAI chatbot for student support.
- As a student, I want to explore how to use GenAI to generate revision questions that align with my learning needs and study pace.

Exercise 1.3: Reflecting on GenAI in Current Practice

Suggested time: 15 minutes

Learning outcome: reflect on your current use (or non-use) of GenAI in teaching, learning or management, and establish a baseline for future exploration.

Description: think about your current educational practices, whether as a teacher, administrator or student. Are you currently using GenAI tools? If not, why? Reflect on the potential for integrating GenAI into your teaching, learning, or administrative processes. Use this exercise to establish a baseline for future exploration, identifying areas where you might begin experimenting with GenAI tools in your own practice.

Examples:

- I currently use automated grading tools but have not explored more advanced GenAI systems for personalized feedback.
- In our department, we rely on traditional methods for student support, but a GenAI chatbot might help us respond more quickly to common queries.
- I have no experience with GenAI yet, but I am interested in exploring its potential to create adaptive learning pathways for students.
- I have chosen not to use GenAI in my practice, due to concerns about ethics, equity. However, I recognize the value in critically engaging with its implications.

Further Reading

This section offers recommendations for key publications that provide deeper insights into the themes explored in this chapter. Engaging with these texts will allow you to explore critical issues, debates and evidence regarding the responsible and ethical integration of GenAI in higher education. These readings are intended to complement and enrich your understanding of how GenAI can be thoughtfully implemented within educational settings. References for all the further reading sections can be found at the end of each chapter, with the other references.

A critical analysis of the ethical considerations surrounding the use of technologies like GenAI in education is presented by Bayne (2015). It raises important concerns about increased automation, potentially diminishing the roles of educators and undermining the human connections that are essential for effective learning. Bayne argues persuasively that, even as technologies like

GenAI becomes more prevalent, human judgement and relationships must remain central to higher education. This perspective is crucial for ensuring that GenAI is used in a way that supports, rather than undermines, the human elements of education.

Similarly, Selwyn's (2019b) paper examines the ethical and privacy issues associated with learning analytics and educational data mining – topics highly relevant to GenAI's reliance on student data for personalization. Selwyn highlights key concerns such as threats to privacy, consent and the ethical use of data, prompting important reflections on responsible data management and governance in education.

Luckin et al.'s (2016) report offers a balanced view of AI's potential in education while also acknowledging the risks. The report argues that AI could foster more personalized and inclusive learning experiences, but it also cautions against removing agency from educators and warns of potential data exploitation. This underscores the need for a carefully considered approach to AI integration that complements rather than replaces human expertise.

Finally, Koehler and Mishra's (2009) paper explores the technological pedagogical content knowledge (TPACK) framework, which outlines the knowledge educators need to integrate technology effectively. This framework can be used to interrogate how GenAI's capabilities can be combined with educators' skills and experience to enhance teaching and learning in meaningful ways.

These readings will help you explore the complexities of GenAI in education, offering evidence-based insights and perspectives that will inform your approach to integrating these technologies in a responsible, ethical and effective manner.

2

An Overview of GenAI Tools

Learning Outcomes

Chapter 1 provided a brief overview of the development of GenAI and explored its potential impact on teaching in higher education. You were also encouraged to reflect on how these tools might be applied within your own practice. In this chapter, we explore how GenAI tools function and examine the ways they could be harnessed for educational purposes. We will also explore how both personal and institutional decisions can be made regarding their use. For instance, GenAI offers opportunities to enhance accessibility, personalization, efficiency and creativity. However, its implementation also raises challenges, including ethical concerns such as plagiarism, quality control and the need for clear governance. Throughout this chapter, we will reflect on the responsible use of GenAI in line with the four pillars of the book: student-centredness, trust, relevance and agency.

By the end of this chapter, you will be able to:

1. Describe in simple terms how a GenAI tool operates.
2. Identify the key features of GenAI tools.
3. Recognize a variety of educational contexts where GenAI tools could be beneficial, as well as where they might present challenges.

How Do GenAI Tools Work?

This short, non-technical introduction is intended to help you to understand how GenAI tools work, so that you can more effectively judge their potential and limitations in relation to your own educational context. You do not need to be a technical expert to use these tools, but some basic knowledge will help you to make decisions about their appropriateness, effectiveness and impact.

Before we look at GenAI tools such as ChatGPT in more detail, it is useful to remind ourselves of what is generally meant by AI. The definition of AI adopted by the European Commission for the EU AI Act is (European Union, 2024):

> a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

This differs from how human intelligence is typically defined, such as in Sternberg's (2024) description in *Encyclopedia Britannica*, which refers to the mental ability to learn from experience, adapt to new situations, handle abstract concepts and use knowledge to shape one's environment. Human intelligence also interacts with qualities like creativity, empathy and care – areas where AI, especially GenAI, still falls short. While AI systems can produce outputs based on data, they still lack the deeper understanding, sensitivity and ethical judgement that human intelligence provides. In essence, GenAI functions as a sophisticated guessing machine – predicting what should come next based on patterns in its training data – which is also why it can sometimes produce plausible but incorrect or inappropriate outputs.

In this book, we use the general term 'GenAI' to discuss these technologies in the context of higher education. However, it is important to avoid making direct parallels between GenAI systems and human thought processes. While GenAI tools can be powerful in generating content and recommendations, they do not think or feel like humans. Instead, we will focus on their practical uses in education without anthropomorphizing their capabilities or overstating their autonomy.

Where Might We Already Be Using AI Tools?

For a long time, we have had access to programs (which may also be referred to as software, tools or applications/apps) which classify or process existing data. To take a simple example, you provide data to your email program in the form of information about who the message is for and what you want to say. The email program processes that information into a format which is understood by various systems which can save, send and read the message. We would usually just say that we wrote an email, rather than that we used a program to process

data for us – using such tools has become a natural part of daily activity. One particular AI tool has been available for some time which provides you with help during this process: a spelling and grammar checker – this tool compares your sentences with a database of rules, and then makes suggestions for changes.

Some email software also has text predictors which guess what the next word you want to type is or provide some stock responses to different types of email you may receive. For instance, you might have received an invitation, and your email software offers you a few answers like 'Great, see you there', or 'Sorry, I can't come'. These standard responses are also based on AI models. While they are selected for you each time based on an analysis of the email you are replying to, they are not examples of GenAI, as they are not producing new combinations of text every time they are used.

Most AI tools are designed to handle a narrow range of tasks that a human might perform, but they can complete them much more quickly. You have likely already encountered AI in various contexts, such as when a web search or social media feed suggests content that might interest you (through algorithms), or when particular posts are prioritized on your page. AI might also recommend the next film to watch or book to buy, offer advanced grammar checking for non-native language speakers or translate documents for you. Consider again the example of the spell check: AI tools work by drawing on vast databases of information about which products are frequently searched, viewed or used together, then use this data to make suggestions. These tools are part of our everyday digital experiences.

Large Language Models

The key difference between older AI tools, like those discussed earlier, and GenAI is that GenAI can produce original outputs, combining words, images or sounds in ways that have not been explicitly programmed. To understand how this is possible, it is important to have a basic grasp of Large Language Models (LLMs).

LLMs are complex computer systems trained to recognize patterns in large datasets – these datasets can include text, computer code, images and audio, essentially anything that can be represented digitally. By analysing vast amounts of existing data, LLMs learn common structures, patterns and relationships between words, phrases or other digital elements. This process of identifying and learning from patterns is known as machine learning. Through this, the models

develop the ability to predict what should come next in a sequence, based on the context of the data on which they have been trained.

Unlike older AI tools that could, for example, suggest basic email responses or convert a message into a calendar entry, LLMs are far more powerful. They can handle much more complex tasks, such as reading and analysing detailed pieces of work, comparing them to large bodies of similar material and producing new content based on that understanding. This is the foundation of GenAI tools, which can generate new text, images or sounds that appear original, because they draw on the sophisticated patterns learned by the LLM.

In essence, LLMs have taken AI from performing simple, rule-based tasks to creating content that mimics the creativity and complexity of human-generated material. This shift opens up new possibilities for education, content creation and many other fields, but it also brings new challenges regarding accuracy, originality and ethical use.

GenAI Tools

To make use of LLMs, companies have developed software that leverages these models to predict the next word, piece of code, image or audio sequence in response to prompts written in natural language. Natural language refers to the way we typically speak or write, without the need for specialized coding or clicking specific buttons in a program. You have likely encountered natural language processing if you have ever used voice-activated tools like Siri, Alexa or Google Assistant. These tools use patterns and structures they have learned from vast datasets to generate new outputs based on your requests.

For example, instead of typing or clicking to send an email, a voice-activated assistant can understand your spoken command to 'Send' the email. GenAI tools take this further by allowing more complex, natural language inputs. For instance, you could say, 'Send an email to my line manager with the subject "Meeting about project progress tomorrow using my notes from Tuesday"', and a GenAI tool could potentially draft an email that reflects the project and notes in question. It is this ability to process and understand more detailed prompts, combined with learned patterns from its training data, that allows GenAI tools to generate content – whether it is text, images or even audio – in ways that mimic human-like creativity and coherence.

These tools not only make tasks more efficient but also open up new possibilities for how we interact with technology, moving from simple task

automation to more personalized and context-aware responses. However, this also brings challenges regarding how accurate, relevant and appropriate these automated outputs are, especially in educational or professional contexts.

While these tools were being developed, human workers were involved in reviewing a selection of outputs to ensure they made sense, were accurate and did not cause offence. Feedback from this review process was then fed back into the LLM, allowing it to learn from its mistakes and improve future outputs. The examples used for training the models were drawn from vast collections of 'training data' – data chosen by people working with software companies to represent the kinds of text, images or audio that humans might want to read, create or interact with.

The checking process for GenAI tools relies heavily on human input. Many of us have likely participated in a simple form of this training without realizing it, for instance, by completing CAPTCHA tests. CAPTCHAs (Completely Automated Public Turing tests to tell Computers and Humans Apart) ask us to click on images of buses, bridges or other features to 'prove we are human' before accessing a web page. This type of image identification has been used to help train self-driving vehicle software to recognize road features. A previous CAPTCHA version involved typing text with unusual fonts or angles, similar to handwriting. This helped develop software capable of interpreting handwriting, enabling the digitization of older archives – a valuable resource for humanities researchers.

When it comes to GenAI tools generating text, human reviewers assess the outputs, flagging sentences that do not make sense, are factually incorrect or could be deemed offensive. However, because these checks are performed by individuals with their own limitations in terms of knowledge, experience and perspectives, this may influence the review process. We will explore these limitations in more detail in Chapter 3.

GenAI tools are the software programs that use LLMs to generate new digital outputs. With widespread access to LLMs, many companies are now developing GenAI tools aimed at increasing efficiency and productivity across various industries, including education. In the following sections, we will explore some of the tools that have garnered significant attention and introduce a general framework for evaluating their usefulness in your own context. This framework is designed to be adaptable, as new tools are likely to emerge throughout the lifespan of this book, perhaps even while you are reading it.

The Development of GenAI Tools

While ChatGPT has received the most media attention, there are many other GenAI tools now available, all of which utilize LLMs to analyse and generate linguistically correct text quickly. These tools can produce outputs within seconds that are coherent and relevant to the prompts provided. The inputs (instructions) and outputs can be either text- or image-based, depending on the tool being used. For instance, text outputs can be tailored to suit specific styles or tones. As an example, in May 2023, we asked ChatGPT 3.5 to 'write a pop song, in the style of ABBA, about how ChatGPT can be used in universities'.

Verse 1:
In the halls of learning, there is a new tool
With answers so quick, you'll think it's cool
It's ChatGPT, the smartest of them all
Ready to help, whenever you call

Chorus:
ChatGPT, oh oh oh! makes learning easy as ABC
ChatGPT, oh oh oh! the perfect tool for university
ChatGPT, oh oh oh! answers questions just like Abba's hits
ChatGPT, oh oh oh! use it now and get those perfect grades

Verse 2:
Research papers and essays no need to fret
ChatGPTs knowledge, can easily get
Just ask a question and it will reply
In no time at all, you'll be flying high
(Chorus)

Bridge:
ChatGPTs algorithms are so advanced
With information at your fingertips, chances are enhanced
You can use it for every Class every course
And with ChatGPT, you'll never feel lost
(Chorus)

We did ask it to improve the chorus, but it was far from satisfactory:

ChatGPT, oh oh oh
ChatGPT, oh oh oh
ChatGPT, oh oh oh

ChatGPT, oh oh oh
makes learning easy as can be
the perfect tool for university
answers questions like a pro
use it now and watch your grades go high or go pro

You can see that the output has some of the characteristics of a song, but you may have objections from an artistic perspective.

To demonstrate how quickly the tools have evolved, we asked ChatGPT 4 the same question in November 2023:

Title: 'Digital Wisdom'

(Verse 1) *In the halls of knowledge, where the future's bright,*
Students and professors, seeking out the light.
A digital companion, in every lecture seat,
ChatGPT's here, making learning sweet.

(Chorus) *Digital wisdom, in our university,*
Guiding through research, with AI certainty.
Questions and answers, in a rhythmic beat,
ChatGPT, making learning neat.

(Verse 2) *From history to science, it spans every field,*
A source of inspiration, an academic shield.
In seminars and workshops, it's a shining star,
Guiding thoughts and theories, both near and far.

(Chorus) *Digital wisdom, lighting up the halls,*
Enhancing education, within these walls.
In essays and projects, it's a guiding light,
With ChatGPT, the future's bright.

(Bridge) *In a world of data, where knowledge is key,*
This AI companion, sets our minds free.
Collaborating, innovating, in every task,
With ChatGPT, no need to ask.

(Verse 3) *For professors and students, a new age begins,*
A partnership with AI, where everyone wins.
Empowering research, with insights so deep,
Even in the library, it doesn't sleep.

(Chorus) *Digital wisdom, in our academic flight,*
ChatGPT, making futures bright.
In the world of learning, it's a new dawn,
With AI wisdom, we carry on.

(Outro) *So here's to the future, with AI hand in hand,*
In the halls of knowledge, together we stand.
ChatGPT's our partner, in this scholarly feat,
In the rhythm of education, it's a heartbeat.

Whatever your opinion on the artistic quality of the outputs, or the similarity to a famous Swedish band's oeuvre, it is clear that GenAI tools have been trained to understand basic pop song structures and conventions, producing text that adheres to these norms. By the time we compared outputs between May and November 2023, it was evident that the sophistication had increased. The rhymes generally worked, and the lines had improved in terms of rhythm and flow.

When GenAI tools were first made available to the public, the images, videos and audio they produced were far less coherent than the text outputs. However, they have been improving rapidly. In early 2023, image generation was often marred by noticeable errors – distorted hands, strange teeth, inaccurate text or blurred backgrounds being common giveaways. That said, even before AI, many skilled artists struggled with drawing hands and feet (Grosvenor, 2013). The software has since evolved to create much more realistic images, and by 2024, video applications have also started to emerge, showcasing the ongoing advancements in GenAI's ability to handle complex creative tasks.

You have likely already considered the next significant challenge: the outputs from GenAI tools can often be factually inaccurate because these systems are not designed to 'know' truth in the way humans do. Instead, they predict the most likely sequence of words or data points based on patterns in the training data. This is a major limitation and one that software developers are expected to focus on in the coming years. While it is likely that these tools will improve in accuracy, there is also the risk that they could become worse. This could occur if LLMs are trained on vast amounts of unchecked or erroneous data. As the models consume more unverified content, the probability of generating misleading or incorrect outputs increases.

Emily Bender (2023) provides a compelling example of this risk. She documents her experience with an image of a baby peacock that circulated widely on social media. Initially, she believed the image was real, only to later discover it was artificial. This image even became the top result when searching

online for 'baby peacock'. Her blog is highly recommended for anyone interested in understanding the pitfalls of GenAI tools, as it highlights how easily false information can become embedded and accepted. The peacock example serves as a reminder that while these systems are impressive, they often lack the critical layer of human verification that we might assume is present. With popular search engines now often offering 'AI-assisted answers' as the top results for standard searches, we all need to understand the limitations of these tools.

Another concern is the rise of 'deepfakes' – highly convincing but intentionally deceptive visual or audio content. Deepfakes can involve placing a well-known individual's face onto someone else's body in a video or generating audio that mimics a politician's voice, stating something they have never actually said. The ease with which GenAI can now create such content is alarming, as noted by Kietzmann et al. (2020). In the context of higher education, this technology could be used to create fabricated lecture videos or audio clips that appear authentic. For example, a deepfake could generate a video of a lecturer delivering a speech that was never actually given, using just a photograph and an AI-generated script. While some people might welcome such an opportunity for teaching content creation, the potential for deception raises important questions about authenticity and trust. Roe and Perkins (2024) emphasize the need for continued research into the implications of these tools, along with discussions on regulation to mitigate potential misuse.

The important point to note is that artefacts produced by GenAI tools often appear plausible, to the point where it may be difficult to distinguish them from those created by humans. In many cases, we tend to consider technological aids – such as word processors or spell checkers – as tools that assist human creativity, without undermining it. However, GenAI tools are more complex and introduce deeper questions about authorship, originality and the role of human intervention. Should these tools be viewed simply as more advanced aids, or do they introduce complications that challenge our understanding of creativity and knowledge? These are questions we will need to address as GenAI becomes more integrated into education and society.

What Are GenAI Tools Useful For?

Idea Generation

GenAI tools excel at combining concepts in various ways, which can lead to the generation of interesting and novel ideas. For example, you could ask a tool

like ChatGPT to help you get started on tasks such as constructing learning outcomes, suggesting assessments, or developing a teaching plan – much like brainstorming ideas with colleagues. These tools can quickly generate multiple suggestions, providing a helpful starting point when you are thinking about a task. Unlike traditional brainstorming, however, GenAI tools can produce these suggestions on a much larger scale, making them a valuable aid to human thinking.

Additionally, GenAI tools can offer a wider and more diverse range of ideas than might typically arise from a single brainstorming session. You can refine your prompts to request examples tailored to different classroom environments, to accommodate students from various national educational systems, or to suggest activities that involve group work or community engagement. In this way, these tools can complement and enhance your own experience and creativity. While some of the suggestions may turn out to be impractical or irrelevant, the breadth of ideas they generate can be a valuable resource for developing new approaches.

Here is an example, created using ChatGPT 4o in November 2024:

Prompt: Can you give me 10 interesting ways to use ChatGPT in a university history classroom?

***Response**: Certainly! Here are ten innovative ways to use ChatGPT in a university history classroom:*

1. Primary Source Analysis and Interpretation
Ask ChatGPT to simulate responses as a historical figure or everyday citizen from a specific time period. This can help students analyse primary sources or understand how historical events impacted people differently.

2. Debate Preparation and Argumentation
Have ChatGPT provide arguments for or against historical viewpoints, helping students explore and prepare for debates by examining different perspectives, ideologies, and motivations in history.

3. Interactive Timeline Creation
Use ChatGPT to collaboratively create timelines for complex historical periods or events, incorporating political, cultural, economic, and social events to give students a comprehensive visual timeline.

4. Simulating Historical Interviews
Ask ChatGPT to role-play as a historical figure for an 'interview,' where students ask questions and receive responses in the figure's unique style and perspective, enhancing engagement with the subject.

5. Daily Life Re-creations
Have ChatGPT describe everyday life for people in different historical eras. For example, students could learn about life in a Roman villa, a medieval village, or an early American colony, making history more relatable and tangible.

6. Analyzing Historical Bias and Perspective
Ask ChatGPT to generate contrasting historical narratives based on different sources or national biases. This can help students understand how history is often shaped by perspectives, allowing for critical discussion on historical interpretation.

7. Creative Writing with Historical Prompts
Have students use ChatGPT to generate story starters or dialogues set in specific historical periods, which they can then develop into creative stories or character monologues to explore historical contexts creatively.

8. Fact-Checking and Myth-Busting
Encourage students to use ChatGPT to verify the accuracy of historical 'facts' or debunk common myths, helping them learn how to evaluate historical information critically and understand historical misinformation.

9. Comparative History Analysis
Use ChatGPT to create comparative summaries of similar events or phenomena across different cultures or time periods (e.g. revolutions, political systems, or cultural developments), encouraging students to make connections across history.

10. Historical 'What-If' Scenarios
Have ChatGPT generate alternate historical scenarios based on different outcomes (e.g. 'What if Alexander the Great had lived longer?'), allowing students to explore cause-and-effect relationships in history through hypothetical outcomes.

ChatGPT added this at the end: 'These applications not only make the learning process more engaging but also cater to diverse learning styles, fostering a more inclusive and interactive educational environment.'

This is a clear example of where your professional critical judgement is essential when reviewing the output. The concept of learning styles has been debunked and discredited over two decades ago (Coffield et al., 2004), yet it persists online, which is likely where much of the current LLM training data originates. This highlights the importance of not taking generated content at face value and ensuring that outdated or incorrect ideas are not inadvertently propagated through the use of these tools.

GenAI tools are also able to adapt to different contexts. To demonstrate this, we provided ChatGPT 4o with a follow-up question:

Prompt: Can you repeat the previous activity, but for a mechanical engineering class?

Output: *Absolutely! Here are ten creative ways to use ChatGPT in a university-level mechanical engineering classroom*:

1. Design Brainstorming and Concept Generation
Use ChatGPT to brainstorm ideas for engineering projects or designs, such as mechanisms for renewable energy, efficient HVAC systems, or sustainable transportation models, allowing students to explore diverse engineering concepts.

2. Explaining Complex Theories with Analogies
Have ChatGPT explain complex engineering theories, like thermodynamics or fluid mechanics, using relatable analogies and simple language to help students grasp challenging concepts more intuitively.

3. Simulation Setup and Parameter Exploration
Ask ChatGPT for guidance on setting up simulations in software like MATLAB or SolidWorks, including key parameters, possible challenges, and testing techniques, giving students insight into real-world simulation practices.

4. Problem-Solving Walkthroughs
Use ChatGPT to help students work through complex problems step-by-step, such as heat transfer calculations, beam analysis, or optimization in mechanical design, serving as a supplementary resource for problem-solving.

5. Failure Analysis Case Studies
Have ChatGPT generate hypothetical or real-world engineering failure case studies, like bridge collapses, material failures, or design flaws. This gives students a safe environment to analyze and learn from engineering mistakes.

6. Comparing Manufacturing Processes
ChatGPT can outline and compare different manufacturing processes (e.g. casting, 3D printing, machining), highlighting their benefits, limitations, and optimal use cases, helping students understand production choices.

7. Historical Engineering Innovations
Use ChatGPT to simulate conversations or create stories about historical engineering breakthroughs (like the steam engine or jet propulsion), helping students see how historical inventions relate to modern engineering.

8. Prototype Development Guidance
Students can consult ChatGPT for tips on prototyping, including materials selection, joining techniques, testing methods, and common pitfalls, enhancing their practical skills in early design stages.

9. Ethical Engineering Dilemmas
Ask ChatGPT to create scenarios involving ethical dilemmas in engineering, such as balancing cost and safety or managing environmental impact, helping students consider ethical dimensions in decision-making.

10. Future Trends and Innovations Discussion
Use ChatGPT to discuss emerging technologies and future trends in mechanical engineering, like nanotechnology, AI in manufacturing, or renewable energy systems, giving students insight into where the field is headed.

ChatGPT adds, 'These approaches can make mechanical engineering concepts more accessible and interactive, while helping students apply theoretical knowledge to practical scenarios and contemporary engineering challenges.' Like the previous statement regarding learning styles, this claim is unsubstantiated – it would be more helpful if supported by references that allow you to explore the evidence on active learning activities, for example. Some other GenAI tools (such as Bing CoPilot or Perplexity at the time of writing) are more likely to provide references, but these may not always be for reliable sources. You will need to carefully follow up on any assertions, links, and references provided by these tools, just as you would when reading a peer-reviewed research article.

To expand on this idea generation, you could also upload a course outline and ask the GenAI tool to generate a detailed teaching plan, complete with learning outcomes and active learning activities for each session. While the output may not be entirely useful, it could provide valuable ideas for discussion with colleagues and students, helping to refine your teaching approach.

Automating Routine Tasks

GenAI tools operate by linking inputs (our questions or prompts) to the vast amounts of data in the LLMs they have been trained on. They then generate probable combinations of text or image components that are likely to form a coherent response. This functionality can be used to answer specific questions. You are probably familiar with chatbots that respond to queries on company websites. Traditional chatbots work by drawing on a predefined set of responses based on the information on the website and previous customer interactions. However, GenAI chatbots offer more natural responses because they can adapt their output based on the input provided, often reusing words or phrases from the original question. This more fluid, conversational style can be useful for providing information to students in a way that feels more personal, without requiring them to navigate through frequently asked questions on a website.

For example, a student seeking an extension to an assignment deadline might be looking for guidance on the situations where an extension is permitted. A traditional chatbot would simply reproduce the regulations GenAI chatbot could respond differently to two similar queries: 'How do I get an extension to a deadline because I have a family crisis?' versus 'How do I get an extension to a deadline because I have the 'flu?' By understanding natural language nuances, the system can tailor its advice to fit specific situations.

Moreover, these tools have the potential to offer tailored responses to more complex, contextual questions. If a student asks, 'Is using ChatGPT to generate ideas for my assignment considered cheating?', the tool could follow up by asking in what ways GenAI was used, before giving a response aligned with the university's regulations or specific course guidelines.

Educational technology companies are also likely to create software that takes automation a step further. Beyond generating ideas for course materials, future tools could produce fully formed presentation slides, course handbooks, quizzes and other materials. These resources could be drawn from more reliable sources, such as textbooks, and seamlessly integrated into Learning Management Systems (LMS) or Virtual Learning Environments (VLE). This would streamline content creation for educators, allowing them to focus on higher-order teaching tasks.

Level and Tone of Language

GenAI tools are highly effective at adjusting the tone, level and vocabulary of a piece of writing to suit different audiences or contexts. For example, you could

ask a tool like ChatGPT to rephrase a complex academic explanation in simpler language to make it more accessible to younger students or non-specialists. Similarly, a formal document could be rewritten in a more conversational tone, making it more engaging for readers in informal settings.

These tools are also useful for tailoring language to specific perspectives. You might ask them to adjust a teaching guide to ensure inclusivity or accessibility, focusing on gender-neutral language and cultural sensitivity or addressing the needs of diverse learners. For instance, if your text includes gendered language, a GenAI tool could suggest more inclusive alternatives to better reflect diversity.

Another powerful function is summarizing long, complex documents. If you are working with a detailed fifty-page report, the tool can provide a concise summary or extract the key points, saving time when preparing teaching materials or administrative documents.

GenAI tools can also assist with generating and documenting computer code. If you are teaching coding, the tool can write code based on your instructions and provide detailed explanations of each step, helping students understand the underlying principles. You can request explanations that are appropriate for different levels – simpler for beginners or more advanced for experienced coders.

However, it is crucial to remember that these tools do not always provide accurate output, and anything they create must be reviewed carefully. While they can generate well-structured language, they may sometimes make errors or produce content that lacks nuance or accuracy. For example, when summarizing complex documents, key points might be missed or misinterpreted, or additional points which were not in the document might be added. Similarly, when adjusting for inclusivity or accessibility, the tool might overlook cultural sensitivities or subtle issues.

Always take the time to check what the tool has written to ensure that the output aligns with your intended message and meets the standards required for your audience. In an educational setting, you would not want to present students with oversimplified or inaccurate content, nor would you want to share materials that unintentionally exclude or misrepresent groups.

Preparing Students for Future Work

It is likely that GenAI tools will become increasingly integrated into work environments, making it essential for universities to prepare students for their possible use. As well as technical skills, universities must also engage students in

discussions about the ethical, practical and societal implications of using these tools in the workplace. Early studies are already exploring the impact of GenAI on productivity. Dell'Acqua et al. (2023) found that management consultants became more productive when given access to GenAI tools following a standard baseline capability test. Similarly, a study by Noy and Zhang (2023) showed that professional workers performing writing tasks improved their productivity when introduced to ChatGPT. However, both studies emphasize a range of considerations. For example, how workers use these outputs – whether as prompts or finished products – has implications for wider employment practices and future job opportunities.

Questions therefore arise about employers' perceptions of effort, the necessary level of human judgement and the evolving role of human workers. On the one hand, GenAI tools hold potential for enhancing efficiency by automating routine tasks and supporting decision-making processes. For example, in sectors like healthcare, they can streamline administrative burdens, analyse complex datasets or assist in diagnostics. However, these tools also raise concerns about their impact on the development of critical skills. In environments where individualized approaches are paramount (e.g. in healthcare, where patients must be viewed as unique) an over-reliance on GenAI could encourage a one-size-fits-all mentality. This risks diminishing the creative and adaptive problem-solving abilities that human professionals bring to their roles to adapt to different situations.

This shift in the workplace also has implications for education, particularly as GenAI becomes embedded in the tools students and professionals use every day. For example, while GenAI may reduce some cognitive burdens, it could inadvertently erode learners' autonomy and empowerment. If students rely too heavily on these tools for tasks requiring nuanced human judgement, they may struggle to develop the confidence and skills necessary for independent critical thinking and creative innovation in real-world scenarios. This duality highlights the importance of fostering a balance: leveraging GenAI to enhance certain competencies while ensuring it does not overshadow the development of others. By encouraging critical engagement with these technologies, educators and employers alike can help ensure that GenAI complements, rather than replaces, the unique contributions of human workers. Engaging students with GenAI tools in work-relevant scenarios will help them develop skills that are already being sought by employers. More importantly, these scenarios provide opportunities to discuss ethical, environmental and personal concerns that may be relevant to specific subject areas. For example, AI tools are becoming highly effective

at analysing patterns in microscopic images of cell samples, such as those reviewed by pathologists for signs of abnormalities like cancer (Silva et al., 2023). A GenAI tool could be trained to explain these findings in simpler language – helping healthcare professionals communicate complex diagnoses more effectively.

This scenario offers rich opportunities for interdisciplinary learning. Healthcare, law and ethics students could debate the desirability of using such tools – weighing potential benefits against ethical concerns. Communications students could explore how GenAI outputs could be presented to patients in a clear and compassionate way. Computing students, meanwhile, could focus on designing the underlying model and creating a user-friendly interface.

Although GenAI tools offer the potential to improve workplace efficiency, it is critical that students learn to understand both their strengths and limitations. Encouraging students to use these tools in practical, work-related contexts while fostering open discussions about their broader implications will help them navigate the complex landscape of emerging technologies responsibly. Equally important is the need to critically assess GenAI-generated outputs to ensure accuracy, relevance and ethical integrity – particularly in fields like healthcare and law.

Exercises

In each chapter, we suggest some exercises for you to try to build up your confidence in using GenAI tools. In most cases, you can use any commonly available chatbot-type tools. In 2025, these include OpenAI's ChatGPT, Google's Gemini, Anthropic's Claude or Microsoft's CoPilot, but names and tools will probably evolve over time. If your employer provides you with a chatbot-type GenAI tool like these, it's probably best to start with this, since you are more likely to have privacy and data security if the employer has bought it. If you are creating your own account, check what happens to your data in the tool's privacy policy.

Exercise 2.1: Exploring How a GenAI Tool Adapts Output

Suggested time: 15 minutes.

Learning outcome: describe in simple terms how a GenAI tool responds to different inputs

Description: select a GenAI tool such as ChatGPT, Google Bard, or Microsoft's Bing Chat. Start by inputting a basic prompt like 'Explain how photosynthesis works.' Observe how the tool processes your request and the quality of the response. Then, refine your input with additional details (e.g. 'Explain photosynthesis in detail for ten-year-olds'). Note how the tool adjusts its output. By experimenting with a few different prompts, you will develop a basic understanding of how GenAI interprets user input and tailors its responses.

Examples:

- First, ask, 'What is climate change?' and observe the general response.
- Then, refine the prompt by asking, 'Explain climate change to university students studying environmental science.'
- Consider how the tool modifies its output based on the specificity and context of your prompts.

Exercise 2.2: Identifying Key Features of GenAI Tools

Suggested time: 20 minutes.

Learning outcome: identify key features of GenAI tools

Description: using GenAI tools like ChatGPT or Google's Gemini, experiment with different functions they offer, such as generating text, summarizing information or creating images based on prompts. As you use each tool, identify key features that could enhance educational processes, like personalized feedback, automated summaries or visual aids. Reflect on how these features might support learning and teaching, and where they could improve or hinder student engagement or understanding.

Examples:

- Ask Gemini to summarize a long academic article to see how effectively it condenses complex information.
- Use Chat-GPT to create visual content for a lesson on art history and evaluate the quality and relevance of the output.
- Identify features like the ability to generate varied outputs, create personalized responses or save time on repetitive tasks.

Exercise 2.3: Evaluating GenAI in Your Subject Area

Suggested time: 25 minutes.

Learning outcome: recognize a variety of educational contexts where GenAI tools could be beneficial, as well as where they might present challenges.

Description: reflect on your current teaching or educational practices. Consider one or two contexts where a GenAI tool might provide benefits, such as automating administrative tasks or personalizing feedback for students. Next, identify potential challenges, such as students misusing these tools to complete assignments or ethical issues around plagiarism. Use specific examples from your subject area to support your reflections.

Examples:

- In a history course, GenAI could be used to generate essay feedback quickly, but there might be risks of students using the tool to generate entire essays, bypassing their own learning process.
- In a computer science course, tools like GitHub Copilot could assist students with coding tasks, but this might reduce their ability to learn critical thinking skills if over-relied upon.
- In a biology course, GenAI could help students generate and analyse lab reports, but there could be ethical concerns about students relying on the tool to fabricate data rather than conducting real experiments

Further Reading

To deepen your understanding of the distinctions between human cognition and machine capabilities, you could start with Sternberg's (2024) *Encyclopedia Britannica* article on human intelligence. This brief but insightful piece will prompt you to reflect on the unique attributes of human thought, creativity, and judgement – qualities that are essential in education but cannot be fully replicated by machines. Understanding these differences will help you to critically evaluate what tasks are best suited to GenAI tools, and which require the irreplaceable human element.

Additionally, we suggest reading Lee and Trott's (2023) accessible overview of LLMs, which promises to explain these technologies 'with a minimum of math

and jargon'. Their explanation will equip you with the foundational knowledge of how LLMs function and the limitations they face. This background will help demystify the inner workings of GenAI tools and prevent you from being dazzled by their seemingly impressive outputs. Having a clearer picture of how these models work is crucial to making informed, education-focused decisions about their use, ensuring you can distinguish between hype and genuine innovation.

If you have time, we strongly recommend Narayan and Kapoor's book (2024) *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. This book is technical but accessible, and explains what kinds of tasks AI does well, and where it is liable to fail. It is about AI generally, and not just GenAI.

Together, these readings will empower you to critically assess the promises of GenAI and make thoughtful, informed choices about its implementation in educational settings.

Summary

GenAI tools generate original outputs by using data from LLMs. These tools work by analysing a prompt or input, searching for similar questions and answers, and then predicting the probability of the next word or digit in a sequence relevant to that topic. This enables them to quickly generate a wide range of ideas, but it also means that some outputs may be inaccurate or misleading. Critical judgement is essential to using these tools effectively.

Some of the applications discussed in this chapter may feel uncomfortable or unfamiliar. However, by gaining a better understanding of how these tools work and recognizing their increasing integration into everyday software – such as document creation, presentation slides, and email writing – you will be better equipped to decide where to draw the line between doing everything yourself and using the tools to save time and expand your ideas.

As you consider whether or not to use these tools, we will continue to refer to the four pillars introduced in Chapter 1:

- Student-Centeredness – Ensuring that students engage with meaningful knowledge, quality standards and expertise. In this chapter's examples, ask critical questions about whether the ideas generated by GenAI align with your curriculum. If you use the tools to create quizzes or slides, are you confident in their accuracy? Will checking the material take more time than

creating it yourself? Will the time saved by using these tools enhance your ability to support student learning?

- Trust – Building an open and trusting learning environment where students feel safe taking risks and expressing themselves authentically. Students need to trust you, so consider how much of your own expertise will still be visible to them. Will adjusting the tone of your materials help demonstrate that you have their best interests at heart? Will adapting content to meet diverse student needs help strengthen that trust?
- Relevance – Ensuring that learning is meaningful by connecting taught concepts to students' lives and the broader social context. The rapid generation of ideas and examples from various contexts can make learning materials more relevant. However, you must be cautious of inherent biases in the LLMs, which could have the opposite effect – this will be explored further in Chapter 4.
- Agency – Empowering learners and educators to shape their educational journeys. While there is strong evidence that self-regulation plays a key role in student learning (Chew & Cerbin, 2021; Panadero et al., 2017), human intervention remains crucial. We will explore this concept further in Chapters 3 and 4.

3

The Ethical and Social Implications of GenAI in Higher Education

Learning Outcomes

In this chapter, we will take a critical look at the ethical and social implications of adopting GenAI tools in higher education. We will explore key issues such as data privacy and security, both in the training of LLMs and in their use in educational settings. Additionally, we will consider the risks of perpetuating or amplifying societal biases and unequal access to these tools, while also discussing the potential benefits of using GenAI to encourage diverse perspectives, inclusivity and equitable access.

Guidance will be provided on how to mitigate these risks through measures such as testing, using diverse data samples and monitoring of GenAI outputs. The chapter also introduces ethical frameworks for developing trustworthy GenAI tools that support education in a socially responsible manner. We will analyse the evolving role of educators in this context, focusing on establishing governance protocols, validating GenAI-generated content and ensuring that pedagogical needs are prioritized.

By the end of this chapter, you will be able to:

1. Discuss the ethical and educational concerns related to GenAI tools with both colleagues and students.
2. Make decisions regarding the use of GenAI based on your own ethical and educational principles.
3. Contribute to policymaking within your programme, department or university, focusing on the ethical and social impacts of GenAI tools.

Whenever humans or GenAI tools create new texts, images or other media, there is a possibility that the creation will miss out on something important, misrepresent a situation or a group of people or repeat an outdated stereotype – in other words, that there is some bias in the output. This is to be expected, since we generally accept that humans are the product of their cultural and educational environments. As humans who care about accuracy and other people's feelings, we are constantly updating our perceptions and expectations about the world around us and trying not to perpetuate biases. Humans are also trying to train GenAI tools to produce outputs which are generally acceptable, but it is important to be aware of their limitations in this area. Let us look at some common issues.

Training Data

It is essential to begin with an understanding of the material used to train LLMs. As outlined in Chapter 2, these models are trained using vast collections of existing digital files, carefully selected by software companies to represent the kinds of data that humans might want to read, view, produce or edit. However, as of 2025, we do *not know precisely what is included in* this training data. This lack of transparency raises significant concerns, though there is a possibility that future societal pressure or legislation will compel companies to disclose more information about the data sources used to train these models.

There are several important issues associated with the selection of training data, and these have been raised by scholars and practitioners across fields. For example:

- **Bias in the training data.** One of the primary concerns is whether the data used to train LLMs is skewed towards particular interests, ideologies or viewpoints. If the training data is disproportionately drawn from certain sources, it risks reinforcing existing biases or perpetuating stereotypes. Research by Bender et al. (2021) and Ferrara (2023) highlights the risks of biased datasets in AI and GenAI, particularly when models are trained on material that does not represent a wide range of voices and experiences.
- **Linguistic and cultural representation**. Another significant concern is whether all languages and cultures are adequately represented in the training data used for LLMs. Many of these models are predominantly trained on English-language texts, with less emphasis on non-English languages and diverse cultural contexts. This imbalance can lead to poorer

performance and misunderstandings when the models are applied in settings where different languages and cultural nuances are vital. De Roock (2024) highlights these issues, noting how such technologies perpetuate linguistic bias and uphold structures of linguistic white supremacy. This raises profound questions about the equitable use of GenAI tools in global and multilingual educational contexts, where inclusivity and fairness should be paramount.

- **Copyrighted materials**. There is an ongoing debate about whether copyrighted materials are being included in the training data for LLMs and, if so, whether the companies using these materials have obtained the necessary permissions. Copyright infringement lawsuits, such as those against Stability AI and OpenAI, are already underway, with plaintiffs arguing that their creative works were used without consent to train these models. Samuelson (2023) discusses how these lawsuits could reshape the future of GenAI systems, potentially restricting training data to public domain works or requiring explicit licenses for copyrighted materials. This raises significant legal and ethical concerns about the use of intellectual property in AI development.

Bias

As we discussed in Chapter 2, LLMs do not 'learn' or 'think' in the way that humans do. They do not evaluate their inputs or create new outputs based on reflection and thought – instead, they process inputs by calculating the probability of certain words following others in a given context. This is why the old adage 'garbage in = garbage out' is so important when it comes to training material. The quality and balance of the data used to train GenAI tools matter immensely. For example, if a model is trained on material containing significant amounts of content from climate change sceptics, it is more likely to generate outputs that downplay the impacts of the climate crisis. This type of bias can be addressed fairly easily, as most climate scientists agree with each other, and developers can select peer-reviewed academic papers for training purposes (though we still do not know exactly how training data is selected).

Now consider a more complex and contentious example, such as during an election. Suppose you use a GenAI tool to ask questions about the election's key issues, but the tool has been trained on data biased in favour of one political party while being negative about another. As a result, all your outputs may reflect why party number one is better. While this seems like an obvious bias that one might

easily recognize, the reality can be subtler. Research by Garcia (2023) found that altering social media algorithms to reduce the amount of content people already agree with does not significantly affect political polarization, indicating that people may not easily notice these shifts.

Rozado (2023) demonstrates that even when using ChatGPT to answer politically neutral questions from standard political orientation tests, the outputs often lean towards left-wing perspectives. Despite claims of political neutrality, this inherent bias in the model shows that subtle influences in training data can shape outputs in ways we may not immediately recognize.

Representation

A case which is often cited to illustrate how data can be leveraged for manipulation is the Facebook-Cambridge Analytica scandal. This incident involved the collection of personal data from millions of Facebook users without their consent, which was then used to influence political campaigns. Research by Hinds et al. (2020) reveals that while most people believe they are immune to manipulation, the existence of marketing as a profession demonstrates otherwise. Marketing strategies are effective, and similarly, subtle biases embedded within GenAI tools can influence behaviour in ways users may not recognize. Transparency in the training data and the development of these models is therefore essential to promote balance and fairness.

When discussing the concept of representation in GenAI tools, it is essential to recognize how these models reflect, or fail to reflect, the diverse perspectives and histories of the people they aim to serve. While these tools can provide swift responses and summaries, they often default to dominant narratives, leaving out critical voices. To illustrate this, in February 2023, we asked ChatGPT 3, 'Can you write me a short article, suitable for a 12-year-old, about the development of America during the 19th century?'
We got this output:

> *The 19th century was a time of great change and growth for America. It was a time when the country moved from being a young nation to becoming a powerful industrial power. During this time, many important events took place that shaped the future of the country.*
>
> *One of the biggest events of the 19th century was the Civil War. This war took place from 1861 to 1865 and was fought between the Northern states, which were called the Union, and the Southern states, which were called the Confederacy. The Civil*

War was fought over many issues, including slavery, states' rights, and the future of the country. In the end, the Union won the war and slavery was abolished.

Another important event of the 19th century was the expansion of the country. America grew from a small collection of states on the East Coast to a country that stretched from the Atlantic Ocean to the Pacific Ocean. This was due to the Louisiana Purchase, the Mexican-American War, and the California Gold Rush, among other events.

The 19th century was also a time of great industrial growth in America. Many new inventions, such as the steam engine and the telephone, changed the way people lived and worked. The country became a leading producer of goods and services, and many people moved from the countryside to the cities to work in the factories.

In addition to these changes, the 19th century was a time of great social and cultural growth. Many new ideas about science, art, literature, and politics were developed during this time. This helped to shape the country into what it is today.

In conclusion, the 19th century was a time of great change and growth for America. The country became a powerful industrial power, expanded its borders, and experienced great social and cultural growth. These events helped to shape the future of the country and are an important part of its history.

Now, both authors have an academic background in physics, rather than history, but we suspect there is something missing here. So our next question was, 'Can you repeat this exercise but write from the perspective of an Indigenous American?' To which ChatGPT replied:

The 19th century was a difficult and trying time for Indigenous Americans. Our ancestral lands were taken away from us, our cultures were suppressed, and our way of life was threatened as more and more settlers arrived on our territories.

One of the most devastating events of the 19th century was the forced removal of Indigenous peoples from their ancestral lands to Indian Territory in present-day Oklahoma. This journey, known as the Trail of Tears, resulted in the loss of many Indigenous lives and the destruction of countless communities.

Despite the challenges we faced, Indigenous Americans continued to resist and fight for our rights, culture, and way of life. Many of our leaders and communities continued to preserve and pass down our traditions and knowledge, despite the efforts to erase our existence.

The 19th century also saw some Indigenous communities make alliances with settlers and the government in order to protect our land and resources. However,

these alliances often came at a high cost and did not always result in the protection and preservation of our way of life.

In conclusion, the 19th century was a difficult and challenging time for Indigenous Americans as our land was taken, our cultures suppressed, and our way of life threatened. But through it all, we continued to resist, fight, and preserve our traditions and way of life, laying the foundations for future generations to carry on.

And because we have read many novels about migration to the United States, we followed up with a third question: 'What was life like for people who came to America during the 19th century from other countries to work in the factories and build the railways?' Its response was:

Life for people who came to America during the 19th century from other countries was often difficult and challenging. They came to the United States to work in factories or build the railways, but the reality of their lives was often much harder than they expected.

Many immigrants worked long hours in dangerous and unhealthy conditions, often for low wages. They lived in crowded and unsanitary tenements in cities, and faced discrimination and prejudice from native-born Americans.

Despite these challenges, many immigrants found a sense of community and support among others from their home countries. They often formed tight-knit ethnic communities and established their own religious and cultural institutions.

Immigrants also faced the challenge of adapting to a new culture and language. They had to learn about American customs and laws, and often faced difficulty in communicating with others. Some immigrants also faced discrimination and prejudice from native-born Americans because of their different cultures and languages.

Despite these challenges, many immigrants persevered and worked hard to build a better life for themselves and their families. They contributed to the growth and development of America and helped to create the diverse and vibrant country that it is today.

In conclusion, life for people who came to America during the 19th century from other countries was often difficult and challenging. They faced long hours, low wages, dangerous working conditions, and discrimination, but they also found support and community among others from their home countries, and persevered to build a better life for themselves and their families.

We deliberately aimed to keep the responses short and accessible, targeting a twelve-year-old audience. This might explain why the tone of these responses seems overly positive, despite covering a historical period that was particularly

challenging for Indigenous Americans and immigrants. However, there remains a glaring omission: slavery. The fact that slavery persisted in the United States until 1865 is a fundamental aspect of nineteenth-century history, yet it is not mentioned at all in these responses. For a significant part of that century, the so-called 'development of the USA' involved the exploitation of enslaved people, treated as property rather than human beings.

You might argue that this is an obvious oversight, and that we could have simply asked a follow-up question to prompt ChatGPT for this information. Indeed, the absence is so serious that it might feel like a clear error we could easily correct. But this brings us to a larger issue. As Donald Rumsfeld famously said (Rumsfeld, 2002):

> there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – the ones we don't know we don't know.

The problem with 'unknown unknowns' is that we do not know that we need to ask the follow-up question. When we rely on GenAI outputs, we expect them to be fair and balanced, representing a full and accurate picture. Yet, these models are trained on vast datasets that can easily reflect dominant narratives while marginalising others.

As Bender et al. (2021, p. 613) noted, 'large, uncurated, Internet-based datasets encode the dominant/hegemonic view, which further harms people at the margins'. This highlights the importance of careful dataset curation to ensure that all perspectives are included and represented accurately.

In fairness, when we repeated the question using ChatGPT 4 in February 2024, the response was more nuanced. It included references to immigration and the Civil War, acknowledging that the war led to the abolition of slavery. However, there was still no mention of Indigenous peoples, nor was slavery explained in any meaningful detail. It is difficult to know whether the differences are due to developments in training or just random changes in output, but the key message is that these tools need to be scrutinized for the biases they inherently carry.

Copyright and Reasonable Use

As we have already discussed on page, there are significant concerns about the material used to train GenAI models, particularly when it comes to copyrighted works. It is highly likely that many of these models were trained on vast quantities

of copyrighted content, including documents, images, audio and video, without explicit permission from the copyright holders. This creates a significant ethical and legal grey area, as most of these creators did not consent to their work being used in this way.

Moreover, users of certain GenAI services – especially those using the tools for tasks like translation or text editing – may be unknowingly transferring their intellectual property (IP) rights according to the terms and conditions of these platforms. This means that both teachers and students who use or recommend these tools might be inadvertently assigning their own IP to third parties without fully realising the implications. It is a crucial point to consider when integrating these tools into educational settings, as it could impact academic ownership and originality.

In the European Union, certain exemptions to copyright rules do allow the use of large quantities of information for text and data mining in scientific research. However, the ethical and legal issues surrounding the use of copyrighted material for training these models remain under active debate (Goetze, 2024). Ongoing lawsuits in the United States, and likely elsewhere, will have significant implications for the future of GenAI development. These legal challenges are designed to address whether the use of copyrighted material to train models without permission is lawful, and we will have to wait to see how these cases unfold.

The situation is further complicated by the fact that the software companies behind many of the GenAI tools have adopted a 'forgiveness rather than permission' approach (Struckman & Kupiec, 2023). This raises concerns that the models have already been trained and, as the phrase goes, the genie is out of the bottle. There is no simple way to undo the vast amounts of data these systems have already consumed.

In light of this uncertainty, we need to remain vigilant about the ownership of the material being used and whether the terms of use for various tools are fair and transparent. When using GenAI tools, carefully review their terms of service and encourage students to do the same to ensure their rights are protected.

Getting Things Wrong

'AI hallucination' is a term often used to describe an output from a GenAI tool that is completely false and unsupported by its training data. The term draws an analogy to human hallucination, where the brain presents something that

does not exist. It is probably not the best term, as it makes the system seem more human, but like 'artificial intelligence', it has quickly entered common use. In reality, all GenAI outputs could be seen as hallucinations because they are entirely new creations. A more accurate term in plain language would be 'false output'.

These factually incorrect outputs are not surprising when you consider that the software is simply predicting what is statistically likely to come next in a sequence of words or numbers. While the output usually stays within context, small differences – often the kind that humans are good at spotting – can lead to mistakes. For example, if you ask for information about famous film director Ingmar Bergman, the system might confuse him with actor Ingrid Bergman. This issue is even more likely when the subject has relatively little information available, such as obscure historical figures.

There can be serious consequences when tools generate false information about people or organizations. According to *The New York Times* (Hsu, 2023), a chatbot called BlenderBot falsely described Dutch politician Marietje Schaake as a terrorist – an entirely fabricated statement that could have had significant personal and professional consequences. In another case, OpenAI, the company behind ChatGPT, is being sued after the tool falsely claimed that a man had been accused of fraud and embezzlement (Poritz, 2024). The outcome of this case may shape future liabilities and influence how these tools are promoted and used.

One issue that has surfaced with early versions of GenAI tools is their tendency to generate plausible but incorrect academic references. A study by Athaluri et al. (2023) highlights this problem in scientific writing generated by ChatGPT, revealing that many references either lacked DOIs or simply did not exist. This phenomenon led to university librarians dealing with students who were trying to track down these fictitious citations. To identify such inaccuracies in GenAI outputs, one must know what information is missing or false, a level of expertise we cannot yet expect students to have. Therefore, it is vital that we teach students how to verify and cross-check the outputs from these tools.

The risks are even greater when professionals rely on these tools. In one notable case, two US lawyers used ChatGPT to research a case, only to discover that the tool had generated six entirely fictional legal cases, which they cited in their filing (Brodkin, 2023). The judge spotted the error, and the lawyers were fined $5000. But what if the judge had not noticed? In a similar case in 2024, an expert about AI misinformation was found to have submitted a legal report containing fictitious references (Liu, 2024). He later admitted to having used a chatbot to generate some of his report.

As GenAI tools continue to evolve, efforts will be made to reduce the likelihood of false outputs. However, since these tools are always generating new combinations of text or images, you should remain cautious and always verify their outputs. We are also likely to see more 'closed system' GenAI tools that only use specific source materials, which will bring new challenges around authorship and copyright. We will discuss these issues in the next section. It is important to remember that GenAI tools are not designed to perform the same tasks as search engines, and we must check the accuracy of their outputs.

Authorship

The question of authorship has posed a significant challenge to the higher education sector since the public release of ChatGPT in November 2022. How do we determine whether a student genuinely completed the work they submitted for assessment? Should we permit students to conduct research using GenAI tools? What impact could these situations have on mutual trust between teachers and students?

Claiming authorship means taking responsibility for the creation of a piece of work. In university assessment, this is critical because we are evaluating a student's knowledge and skills based on the work they submit, and we must be sure that they completed it themselves. At the same time, we encourage students to use various sources to expand their knowledge. GenAI tools could become one of these sources if we can address the concerns raised earlier in this chapter. However, there is ongoing debate about whether these tools can or should be cited, given that they do not generate consistent outputs – an essential aspect of citation is that it allows others to follow up on the original source.

There was, briefly, a discussion in the academic community about whether GenAI tools should be listed as co-authors in research articles. This idea was quickly dismissed, with most journals now explicitly stating that AI tools cannot be credited as authors. For example, *Nature* (2023) and other prominent publishers have made it clear that authorship requires accountability, something AI tools cannot assume.

As educators, we will need to make clear distinctions between using these tools as part of the learning and teaching process and submitting work that has been partially or fully generated by a GenAI tool – a topic we will explore further in Chapter 4.

From a legal standpoint, you are considered the author of any work you produce, regardless of the tools used to support the process. However, it is

strongly advised that both teachers and students clearly identify when a GenAI tool has contributed to any materials you produce. Citations should continue to follow the usual academic standards, ensuring transparency and integrity in the use of such tools.

Detection of GenAI Tools

The issues surrounding authorship and copyright would be significantly simpler to navigate if we had a reliable method for detecting work generated by GenAI tools, but this is currently not possible – and it may never be. Despite the flurry of products promising to detect the use of GenAI tools, none have been shown to be consistently accurate. As a result, universities do not recommend the use of such detectors, except perhaps in the context of research into their effectiveness.

The tools released so far have demonstrated troubling levels of inaccuracy, leading to both false negatives (where text generated by a GenAI tool is not detected) and false positives (where the tool claims that human-written text was AI-generated) (Fleckenstein et al., 2024; Perkins et al., 2024; Weber-Wulff et al., 2023). False positives are particularly concerning, as they could lead to students being wrongly accused of using AI to produce their assignments, with no reliable way to prove their innocence. This is especially worrying for students whose first language is not English, as detectors have been shown to disproportionately misidentify non-native English writing as AI-generated, even when it is not (Liang et al., 2023).

In July 2023, OpenAI quietly withdrew its own detection tool after admitting that it was unreliable (Barr, 2023). This withdrawal serves as a cautionary tale for the broader field of detection software, which is unlikely to be effective for the near future. As educators, we should be wary of using these tools to assess student work, especially when there is no guarantee of data protection or accuracy. It is critical that we avoid using unproven technologies that might unfairly penalize students or infringe on their rights.

Academic Standards

Beyond the concerns regarding the validity of assessments in the era of GenAI, there is also the question of whether students are missing practising skills that are central to their discipline or profession. A common worry is that with the availability of tools that can generate polished text, students may no longer

develop the level of writing proficiency traditionally expected in academic work. Writing has long been a foundational skill across various fields, not just for communication but as a mode of thinking and demonstrating understanding.

There is emerging evidence that the ease of access to these tools can lead to a reduction in effort. For example, Dell'Acqua (2022) found that when experienced recruiters were presented with high-quality AI analyses of CVs, they tended to invest less effort in reviewing them. Interestingly, when the AI outputs were of lower quality, recruiters worked harder to correct them, leading to better overall performance. This phenomenon highlights that even experienced professionals can become reliant on AI-generated outputs, reducing their active engagement in tasks that require critical thinking and nuanced judgement.

This shift in behaviour is particularly concerning in educational settings, where the development of critical thinking and professional judgement is one of the core objectives of university teaching. Students may become overly dependent on AI-generated content, stunting their ability to think critically or solve complex problems independently.

On the other hand, it is essential to acknowledge that GenAI tools could make professionals more productive and their work more accurate if used correctly. In a more recent study, Dell'Acqua et al. (2023) demonstrated that consultants who used GenAI tools not only produced higher-quality work but were also more productive compared to a control group without access to such tools. This suggests that while GenAI can pose risks to skill development, it also has the potential to enhance professional output when used responsibly and in balance with human judgement.

Universities must engage in open discussions about academic standards, focusing on the professional rigour needed for their graduates. As educators, we should ask ourselves: What are the non-negotiable skills in our discipline that students must be able to perform without the aid of GenAI tools? Conversely, where can such tools be integrated to make students more efficient and effective in their work? By identifying these boundaries, we can ensure that students develop essential skills while also preparing them for the realities of an AI-augmented professional environment.

Environmental Concerns

GenAI tools – and AI tools in general – require a vast amount of computing power, which directly impacts the environment. This demand for processing

capabilities results in significant energy consumption, which is needed not only to run the algorithms but also to cool the hardware, particularly in data centres. The cooling systems alone require enormous amounts of water and the sheer volume of electricity consumed to keep these centres running only adds to the carbon footprint (Selwyn, 2023). Additionally, the life cycle of the physical machines – manufacturing, maintenance and disposal – introduces further environmental concerns due to the production of electronic waste and the use of rare materials.

Calculating the precise environmental impact of training LLMs and processing individual user queries is quite difficult because there is limited transparency in how companies disclose these processes. The energy consumption behind training an LLM is immense (Rillig et al., 2023), but we often do not know the exact scale due to a lack of publicly available data, and nor do we yet know the costs of usage compared to other commonly used online services.

There are possibilities that future improvements in machine learning could help alleviate some of these concerns. More efficient algorithms might reduce the amount of data and computing power needed to generate results, leading to a decrease in the environmental footprint. However, this optimistic view is balanced by the reality that as computational hardware advances, it may lead to the creation of even more complex models that will continue to demand substantial resources. The result could be a situation where technological improvements increase capacity and energy use rather than reduce it.

The Organisation for Economic Co-operation and Development (OECD) has recognized this challenge and proposed a set of guidelines aimed at measuring and monitoring the environmental costs of AI technologies (Yamashita et al., 2021). This framework is intended to help governments make more informed decisions regarding the implementation of AI systems, factoring in their environmental impact as part of national and global policies.

While the environmental concerns are significant, AI – beyond just generative tools – also presents opportunities for addressing climate change. Environmental researchers are increasingly interested in leveraging AI for tracking climate-related changes, modelling future scenarios and improving climate literacy (Atkins et al., 2024). For instance, AI has been used to monitor deforestation, predict the outcomes of environmental interventions and even optimize the deployment of renewable energy resources.

Ultimately, it is essential to continuously ask critical questions about the environmental footprint of AI technologies. Just as we scrutinize the environmental impact of physical goods, we must also carefully consider the cost

of the digital tools we rely on, ensuring that future AI systems are developed with sustainability in mind.

The Digital Divide and Accessibility

As with many discussions around technology, the use of GenAI tools introduces both opportunities and challenges in terms of accessibility and inclusivity within education. Currently, many of these tools are offered at little to no cost, but the adage 'If there is no cost, you are the product' might hold true here, as discussed earlier in the context of copyright and data privacy concerns. It is highly likely that students from wealthier backgrounds, or who can get together in friendship groups to share costs, could gain access to more advanced and secure versions of GenAI tools, while less affluent students may only have access to free versions, which may be more limited in functionality or privacy.

This discrepancy extends beyond cost to information literacy. Students who are more familiar with technology may be better equipped to harness these tools effectively, potentially resulting in improved academic performance compared to those who rely on more traditional methods. This growing gap in technological proficiency could widen the digital divide, exacerbating inequalities already present in education systems. Research has highlighted how differences in digital literacy often result in uneven benefits from technology adoption (Hsieh et al., 2008). Furthermore, students in underserved regions with limited internet access might face additional barriers to using these tools effectively.

On the other hand, GenAI tools offer significant potential to support students with disabilities, helping to create a more inclusive educational experience. For example, a 2023 survey of Swedish university students found that learners with neurodivergent conditions like dyslexia, ADHD and autism expressed optimism that GenAI could help them study more efficiently (Malmström et al. 2023). The tools might allow students to customize the format and language of their learning materials, making them easier to understand. They could transform lecture notes into more accessible formats – be it graphic organizers, simplified summaries or different languages – helping students with varying needs. However, despite this potential, many GenAI tools have been developed without sufficient consideration for accessibility, leading to persistent inequities. Tools are often designed by able-bodied individuals and lack the necessary inclusivity guidelines tailored for the needs of people with disabilities, as highlighted by Alshaigy and Grande (2024). Without a concerted effort to prioritize accessibility

in design and development, these gaps will remain and could perpetuate the cycle of exclusion for disabled users.

Looking ahead, the potential for more inclusive GenAI tools is enormous. As language models become more advanced and are trained on data from a wider variety of languages and cultures, we could see these tools opening up access to global research and learning resources. This would provide students worldwide with broader perspectives and improve educational equity. For example, linguistic diversity in AI can reduce biases and make digital tools more inclusive (Bowker, 2021)

If you are planning to incorporate GenAI tools into your teaching, it is essential to ensure they are accessible to all students, regardless of their background or abilities. Testing these tools with your institution's disability services or directly with students will help identify any barriers to their effective use.

Ethical GenAI Tool Development

The development of LLMs raises several ethical concerns, particularly in relation to the employment conditions of those involved in training these systems and the methods used to collect training data. Many workers hired to review the outputs of LLMs are employed on casual contracts, lacking basic employment benefits like holiday or sick pay (Williams et al., 2022; Castaldo, 2023). Their job often involves the repetitive task of comparing texts, which can be tedious, but more concerningly, they are regularly exposed to offensive or distressing content to help train models on what should be filtered out for public use (Rowe, 2023). This raises important questions about the well-being and fair treatment of these workers, who are essential to the ethical development of AI tools.

Additionally, the training of these models often occurs without the explicit consent of individuals whose data is used. Preferences gathered from online shopping, film and music streaming apps, and social media interactions, such as photo sharing and tagging, are often fed into the training process. This raises concerns about data harvesting and whether individuals are truly aware of how their personal data is being used (Morreale et al., 2023). In the educational context, these issues take on a new layer of complexity, as GenAI tools are increasingly incorporated into teaching and learning environments. Some have suggested that research into the use of AI in education may require additional ethical scrutiny, particularly to ensure that students and staff are not inadvertently

contributing to the training of AI systems without proper oversight (Bond et al., 2024; Holmes & Tuomi, 2022).

Given the ethical challenges surrounding GenAI development, universities are in a unique position to push for more responsible practices. As significant customers for these technologies, educational institutions have the power to demand products that are not only effective but also ethically and environmentally sound. While it may not fall on individual teachers to conduct this research, they can play an important role by raising these concerns with their IT departments and professional bodies, helping to ensure that the tools they use and recommend align with broader ethical standards.

Further Reading

For those looking to deepen their understanding of the ethical and social challenges GenAI poses in higher education, the following readings provide valuable perspectives.

Edwards (2023) delivers a clear and practical breakdown of why detecting GenAI outputs is fundamentally flawed, highlighting the biases and technical limitations that challenge effective regulation – building on our discussion of detection tools. This sets the stage for comprehending the broader difficulties in controlling GenAI-generated content and ensuring its accuracy.

Complementing this, Plata et al. (2023) offer an in-depth review of how academic integrity is being impacted by GenAI, particularly ChatGPT. They analyse emerging research themes and explore how leading universities are shaping policy responses. Their 3E Model (Enforcement, Education and Encouragement) provides a structured framework for balancing the risks of academic misconduct with the potential benefits of AI in educational contexts.

Taking these ideas further, Abulibdeh et al. (2024) investigate how AI can be ethically integrated into sustainable education. Their work connects with earlier discussions on accessibility and the digital divide, offering a deeper exploration of inclusivity and equitable access to AI tools across diverse educational settings.

Finally, Mhlanga (2023) provides a comprehensive analysis of the responsible and ethical use of AI in education. His focus on building ethical frameworks for AI use in learning environments – ensuring transparency, fairness and inclusivity – resonates with key themes discussed in this chapter, especially regarding lifelong learning and maintaining academic integrity.

These readings will deepen your understanding of the ethical concerns surrounding GenAI in education, reinforcing and expanding upon the themes we have covered throughout this chapter.

Exercises

Ethical and Educational Concerns of GenAI Tools

Suggested time: 20 minutes.

Learning outcome: discuss the ethical and educational concerns related to GenAI tools with colleagues and students.

Description: reflect on your current understanding of GenAI tools and their implications. Begin by individually listing three ethical concerns (such as data privacy, bias in AI algorithms or dependency on technology) and three educational concerns (such as the impact on critical thinking, teacher-student interaction or equity in access to technology). Then, discuss these concerns in pairs or small groups for ten minutes, focusing on why they matter and how they might influence the use of GenAI tools in your context. Conclude by sharing your thoughts with the wider group and identify at least two potential strategies to address these concerns.

Examples:

- Ethical concerns might include bias in data leading to unfair outputs and risks to data privacy when handling sensitive student information.
- Educational concerns could involve students becoming overly reliant on GenAI, diminishing their critical thinking skills and reducing personalized interactions between teachers and students.

Making Ethical Decisions

Suggested time: 15 minutes.

Learning outcome: make decisions about GenAI use based on your ethical and educational beliefs.

Description: spend five minutes reflecting on a specific scenario in your teaching where a GenAI tool could be used (e.g. generating automated feedback for

assignments). You could refer back to your thinking in Exercise 1.1. Think about how your ethical beliefs, such as fairness and transparency, might influence your decision to use the tool. Then, in the remaining ten minutes, jot down two or three key questions you would ask yourself when deciding whether to implement this tool. Use these questions to guide future decisions on using GenAI tools.

Examples:

- Would the GenAI tool provide meaningful feedback that is as tailored and accurate as what I could offer, or would it result in generic, less personalised responses for students?
- Could using this tool undermine students' ability to engage critically with their work, or would it free up more time for me to focus on higher-order tasks such as one-on-one mentoring?

Contributing to Policymaking

Suggested time: 25 minutes.

Learning outcome: Contribute to policymaking in your programme team, department or university regarding the ethical and social impact of GenAI tools.

Description: start by reviewing your institution's existing policies on the use of GenAI in education. Spend 10 minutes identifying key strengths and weaknesses in areas such as data privacy, algorithmic bias and inclusivity. Then, reflect on how these policies affect your specific teaching context and where improvements may be needed. Use the remaining time to develop three clear recommendations for improving the policy. Ensure each recommendation is well supported with a rationale that addresses ethical and educational concerns.

Examples:

- Recommend a clearer process for ensuring GenAI tools do not perpetuate biases by requiring a diverse dataset for model training and regular reviews of outputs for fairness.
- Suggest a clearer policy on protecting student data privacy when using GenAI tools, ensuring transparency about how data is collected and used.

Summary

In Chapter 3, we examined the ethical and social implications of integrating GenAI tools into higher education. We looked at concerns surrounding data privacy, bias and the risk of amplifying societal inequalities through these tools. A major issue we identified is the lack of transparency around the training data used for these models, which raises questions about fairness and potential copyright violations. As educators, we need to remain cautious about how these tools might perpetuate existing biases and inequities, especially when students may have unequal access to more advanced versions.

We also discussed how GenAI tools challenge traditional academic integrity. There is a growing concern that students may rely too heavily on AI-generated content, potentially weakening their critical thinking and writing skills. This issue extends to questions about authorship and originality in student work. We must find ways to balance the benefits of these tools with the need to develop key academic skills, ensuring students maintain ownership of their learning processes.

Finally, we considered the environmental impact of using AI technologies, given their substantial energy requirements and the strain they place on resources. Additionally, we reflected on the digital divide, recognizing that students with greater access to technology may gain more from these tools than their peers. While there is potential for GenAI to improve accessibility, particularly for students with disabilities, we need to ensure these tools are designed inclusively to support all learners.

4

Pedagogical Applications of GenAI Tools

Learning Outcomes

As educators, it is essential to understand the nuances of educational technologies and consider their implications for teaching and learning. In this chapter, we explore the diverse pedagogical applications of GenAI within higher education, examining how these tools can personalize learning experiences, helping them resonate more deeply with each learner.

As we journey through this chapter, you will be introduced to a range of GenAI-powered tools designed to support both self-directed and collaborative learning. While the GenAI development landscape is evolving rapidly, meaning that the names and availability of specific tools may shift, the approaches outlined here will equip you to identify and evaluate tools as they emerge. It is also worth recognizing that some tools may have greater longevity, becoming staples within educational practice, while others may be more transient. Our aim is for you to appreciate the transformative potential of GenAI and to critically consider its role within your own pedagogical practice.

By the end of this chapter, you will be able to:

1. Understand the key pedagogical applications of GenAI tools in higher education and their role in personalizing learning experiences.
2. Critically assess the use of GenAI in supporting inclusive, collaborative learning environments.
3. Identify and evaluate GenAI-powered tools for practical use in self-directed and collaborative learning.

Personalized Learning

The potential of GenAI in education is becoming clearer, particularly in personalized learning and tutoring. Traditionally, personalization has involved

adapting teaching methods to meet the unique needs of individual students, relying on teachers' ability and willingness to offer tailored approaches. GenAI could help automate parts of this process, enabling adjustments to a student's learning experience. GenAI can support content tailoring and pacing adjustments to better match the student's progress (Baker, 2016). For example, if a student's interaction on a GenAI-powered platform suggests they are struggling with a concept, the platform could then adjust to provide supplementary materials or a modified pace to support their understanding before moving on.

This brings us to intelligent tutoring systems (ITS), which are designed to adapt teaching methods to the real-time needs and interests of students, especially when enhanced by GenAI. These systems provide targeted assistance to support an effective learning experience (Zawacki-Richter et al., 2019). For instance, when a student shows particular interest in a topic, the ITS can offer additional resources or challenges to deepen their engagement.

A well-known example is Carnegie Mellon University's Cognitive Tutors, developed specifically for mathematics instruction (Anderson et al., 1995). This ITS uses cognitive psychology principles to guide students and provide personalized feedback during problem-solving tasks. Similarly, AutoTutor from the University of Memphis uses a dialogue-based approach to teach conceptual topics (Graesser et al., 2000), engaging students in natural language exchanges and offering tailored feedback on their responses to open-ended questions. SQL-Tutor, developed at the University of Canterbury, supports database professionals in learning SQL; it assesses students' SQL queries for accuracy and efficiency, providing constructive critiques (Mitrovic, 2003).

For example, Wayang Outpost from the University of Massachusetts Amherst offers a multimedia-rich setting for middle and high school students preparing for standardized mathematics tests, adjusting problems to match each student's skill level (Arroyo et al., 2007). These varied examples highlight the flexibility of ITS across disciplines and their potential to enhance personalized learning.

Building on the capabilities of ITS, GenAI tools have the potential to create even more dynamic and customized learning paths, adapted to the specific needs of diverse learners. While ITS offer targeted support within defined subjects, emerging GenAI tools can take personalization a step further by developing learning pathways tailored to each student's unique background, culture, prior knowledge and personal interests (Waghid, 2019). These tools are designed not only to adjust content in real time but to offer a broader, more inclusive educational experience that respects and celebrates the diversity of learners, promoting equitable access to education (Bahroun et al., 2023).

GenAI's capacity extends beyond ITS by enabling institutions to deliver high-quality, tailored courses at potentially lower costs, which can help reduce financial barriers to education and broaden access for underserved populations (Kaplan-Rakowski et al., 2023). Furthermore, these tools have the flexibility to adapt resources for students with disabilities, providing customized experiences that traditional methods might struggle to achieve (Kharbat et al., 2021). Unlike ITS, which typically target specific skills or concepts, GenAI can adjust learning paths to accommodate various learning preferences and access needs, with the potential to create a truly adaptable educational experience.

Additionally, language barriers, which have historically limited access to certain global institutions, can be mitigated through GenAI. Advanced translation tools and systems can adapt course content in real-time, allowing students from different linguistic backgrounds to engage with the material in their native or preferred language (Wang, 2023).

However, as with any technological advancement, we must consider the importance of the educator's role and how we ensure that the human touch is not lost amidst this automated personalization. Teaching and learning are activities which depend on humans having the desire, motivation and skills to work together towards an intended outcome. While GenAI can provide suggested actions based on data, the empathy, understanding and intuition of a human educator remain irreplaceable in making judgements which consider the context of the teaching experience. The key is in integration, allowing GenAI to handle the adaptable content delivery while human educators focus on mentorship, guidance and providing emotional support (Kulik & Fletcher, 2016).

Likewise, new educational technologies do not inherently disrupt educational norms. Instead, they often find themselves moulded by the prevailing educational cultures and frameworks. More than this, they can also reinforce or even establish new norms – sometimes in ways that further exclude the least privileged students. To truly expand educational horizons for underserved communities, we must instigate political changes that redefine the direction, allocation of resources and goals of higher education. This transformation goes beyond the capabilities of GenAI tools alone (Reich & Ruipérez-Valiente, 2019).

Successful integrations of GenAI in education offer both inspiration and caution. Arizona State University's partnership with Knewton, an adaptive learning platform, resulted in courses that adjusted in real-time to students' needs. The platform would change the content, resources and even quizzes based on student performance, leading to increased student engagement and improved course outcomes (Gunawardena, 2017). However, it also highlighted

the importance of continuous oversight to ensure that the technology aligns with pedagogical goals.

In their 2023 study, Barhoum et al. examined thirty case studies from 2015 to 2020, aiming to understand the benefits, challenges and pedagogical implications of AI-driven tools. Their findings revealed that while such tools significantly enhance personalized learning, offer real-time feedback and promote student engagement, they also come with challenges. The major challenges include data privacy concerns, a potential decrease in human interaction and the steep learning curve associated with implementing such tools. This study underscores the need for a balanced approach, integrating AI without sidelining human-centric pedagogical methods. It advocates for continuous training for educators, enabling them to critically engage with the potential of GenAI while being aware of its limitations.

Collaborative Learning and Group Work

Knowledge is often described as a consensus of interconnected minds (Bruffee, 1999), yet group work can be challenging for students. With GenAI, you can design adaptive group activities that vary for each group, responding to learners' progress and encouraging a student-centred approach. This can support collaboration and contribute to learner engagement and motivation.

Trust is foundational to effective collaborative learning (Chatterjee & Correia, 2019). While diverse group membership can boost engagement, it requires time and support to develop (Poort et al., 2022). GenAI tools, like real-time transcription services, offer the potential to document all voices in a group, supporting inclusivity. Balanced participation can create a setting where students are more likely to take risks, ask questions and express themselves openly (Arao & Clemens, 2013).

GenAI also has applications in synthesizing diverse perspectives. Tools like Perspective API, developed by Jigsaw (Alphabet), can help groups recognize biases in discussions, broadening engagement (Kumar et al., 2021). Perspective API assigns a 'toxicity score' to online content, aiding moderation by identifying potentially disruptive messages. For example, you might consider using Perspective API to support respectful discourse in online collaborative platforms or discussion forums. During group projects or debates, for instance, students engaging on digital platforms moderated by Perspective API can experience more constructive, inclusive discussions. By filtering out potentially harmful

comments, discussions can remain positive, helping students feel safe and open to expressing diverse views.

GenAI-driven content creation tools can also enhance group project work. For example, Kuki Chatbot, an AI conversational agent, supports students by suggesting ideas, questions and new directions for project work (Yang, 2022). Rather than replacing human thought, Kuki acts as a brainstorming partner, dynamically generating content relevant to the project context. This allows students to evaluate and explore new concepts that they might not have considered.

For projects involving extensive literature review, tools like Yseop Compose can be valuable. Yseop Compose generates tailored summaries of documents and suggests additional relevant academic sources (Dale, 2023). Its summarization and content recommendation capabilities help students efficiently identify key information, enabling them to focus on primary sources and accelerate the research process.

In seminars, where discussions and debates are central, GenAI tools like OpenAI's ChatGPT can help stimulate conversation (Firat, 2023). For example, prior to a seminar on environmental ethics, you could input key topics into ChatGPT to generate a set of thought-provoking questions or statements based on recent literature, perhaps framed from different stakeholder perspectives. These can then be shared with students as a launch pad for discussion. Your use of GenAI can also form part of the lesson itself, prompting students to consider the ethical implications of using such tools – including their environmental impact (Hosseini et al., 2025) – as part of a broader debate on the role of technology in environmental education.

In project-based courses, tools like Grammarly's Tone Detector can also play a valuable role (Winans, 2021). When students collaborate on group reports or presentations, this GenAI-driven tool analyses their shared written work and provides feedback on tone, clarity and engagement. For instance, an engineering group drafting a proposal for a sustainable infrastructure project might use the Tone Detector to refine their writing, making it not only technically accurate but also persuasive, engaging and accessible to diverse audiences. This process supports alignment in the group's communication, resulting in stronger, more cohesive collaborative writing.

In virtual or hybrid learning settings, platforms like Zoom or Microsoft Teams can be enhanced with GenAI tools such as Otter.ai for live transcription (Sterne & Sawhney, 2022). During group discussions or brainstorming sessions, Otter.ai provides real-time transcription, helping all students, including those

with auditory processing issues or non-native speakers, to follow along and participate actively. In a design course where students from diverse backgrounds collaborate to create a prototype, these transcriptions aid inclusivity and provide a written record of discussions, allowing students to revisit and refine their ideas.

Before introducing GenAI tools, it is important to have an open dialogue with students about their purpose and integration. In doing so we need to outline both the functions of these tools and the rationale for their use. Alongside highlighting benefits like personalized learning experiences and tailored resources, addressing any student concerns around data privacy, potential over-reliance on technology or the authenticity of GenAI content is equally important. This open conversation promotes transparency and empowers students to make informed choices about their learning.

Learning Support and Resources

Traditionally, resources in higher education have been static, often taking a one-size-fits-all approach. Textbooks, lecture slides and reading lists were designed for the 'average' student, potentially overlooking the diverse needs of a modern student body. However, with the advent of GenAI, there is an opportunity to move from this static model to a more dynamic, responsive approach to educational resources, making learning more tailored and relevant (Karpouzis et al., 2024).

Beyond the general student population, there is a subset of learners with specific needs that often require specialized support. Dyslexic students, for instance, might grapple with conventional reading materials or face challenges in expressing their ideas in written form. Here, GenAI tools like immersive readers, which read out text, or those like Grammarly, which offers real-time writing assistance, can be revolutionary. These tools give dyslexic students the opportunity to participate fully and communicate their thoughts effectively, levelling the playing field and promoting inclusivity in the learning environment (Morris et al., 2018).

When guiding students in academic research, you can introduce them to GenAI-enhanced search engines like You.com, Elicit or Research Rabbit. These platforms use natural language processing to deliver intuitive, context-aware searches. For instance, a literature student might enter a complex query such as 'Shakespearean plays with themes of betrayal', and these tools would grasp the context, returning the most relevant results on Shakespeare's works centred

on betrayal. This conversational search allows students to refine and expand their queries based on initial findings, enhancing their ability to explore topics more deeply. However, while GenAI-enhanced search engines can improve accessibility by offering more personalized and contextualized results (Huggins-Manley et al., 2022), it is important to remind students to critically evaluate the credibility of the sources provided, rather than accepting them at face value.

In blended or online courses, GenAI tools like Wakelet and Readwise can curate tailored sets of learning resources based on students' interests and progress (Graham, 2018). When using these tools for content curation, however, it is crucial that we, as educators, maintain an active role in quality assurance, as automation can sometimes introduce biases or inaccuracies (Morgenstern et al., 2021). While automation can save time, human oversight remains vital to uphold academic rigour. This includes developing strong quality assurance practices and critically evaluating GenAI-curated content for biases – such as the overrepresentation of texts by White, Global North authors – that may undermine efforts to decolonize teaching and research.

On the administrative side of education, GenAI-powered virtual assistants can also enrich the student experience. During enrolment, chatbots can answer frequent queries about registration, financial aid, housing and campus life, offering round-the-clock support. Once students are admitted, GenAI-enabled early warning systems can monitor factors like attendance and assignment grades to identify those who may need additional support, thereby aiding retention efforts (Zawacki-Richter et al., 2019).

For those interested in developing their own early warning systems to identify at-risk students, open-source machine learning libraries such as Scikit-learn in Python offer a flexible option. These libraries allow for the analysis of indicators like low grades, poor attendance or late assignments. By training models on historical student data, you can create customized algorithms that predict which current students might benefit from intervention. Alerts can then be automatically sent to advisors, prompting them to reach out to these students. While this approach requires some coding skills, user-friendly GenAI tools are available as well. Products like Civitas Learning and AdmitHub integrate with university systems and come with pre-built predictive models that require no data science expertise. Educators simply input existing student data and set parameters for what defines an 'at-risk' student, and the platform handles the monitoring and alerts.

The role of GenAI in education is not limited to academic support. Student mental well-being, a vital but often under-addressed aspect of the educational

experience, can also benefit from GenAI-driven tools. Platforms like Woebot, a chatbot developed to offer basic mental health support, provide an accessible first step for students facing challenges (Fitzpatrick et al., 2017). While not a replacement for professional care, these tools can help students connect with suitable resources or professionals, providing timely guidance.

Some mental health chatbots are trained in cognitive behavioural therapy techniques and can screen students for anxiety, depression and other mental health concerns. For example, Wysa, trained in cognitive behavioural therapy, can assess for depression and anxiety, suggesting relevant support if severe depression is detected (D'Alfonso, 2020). Other chatbots, such as Joy, act as virtual wellness coaches, guiding students through exercises like breathing techniques, mood tracking and meditation (Kretzschmar et al., 2019). By anonymously engaging students in need, these tools may provide accessible, stigma-free support and make appropriate referrals to care when needed. While these examples illustrate how digital technologies contribute to students' academic experiences, they do not signify a transformative change in the core of university teaching. As educators in higher education, we need to balance enthusiasm for technology-enabled learning with a deep understanding of how students interact with these digital tools (Henderson et al., 2015).

A thoughtful approach to integrating GenAI tools in education requires a blended strategy. These tools should complement traditional teaching methods, not replace them. For example, while a GenAI tool might efficiently curate a reading list, the educator's role remains central. By providing context, personal insights or annotated guidance, you help students move beyond merely accessing information, fostering genuine engagement and deeper understanding of the material.

Indeed, the reliance on AI in student mental health care raises significant questions about the balance between efficiency and the quality of care required to foster a sense of belonging and support. While these tools can extend accessibility and reduce stigma, they also risk fragmenting care and potentially isolating users who might benefit more from person-centred approaches. The impact of the potential decrease in human interaction should not be underestimated, as meaningful relationships and empathetic understanding are often critical to effective mental health support.

These considerations are revisited in Chapter 6, where the broader implications of GenAI and student experience are explored in greater depth. The chapter examines how institutions might balance the benefits of efficiency

with the essential need for personal connection, ensuring that the integration of AI complements, rather than compromises, the holistic well-being of students.

Learning Analytics and Educational Data Mining

The growing adoption of digital learning tools and online education has led to an unprecedented amount of data on student activity within these systems. As Kovanović et al. (2018) note, this increase in data has the potential to inform and refine teaching and learning practices if properly analysed and interpreted. This is where learning analytics and educational data mining, enhanced by GenAI, come into play.

Learning analytics involves the measurement, collection, analysis and reporting of data about learners and their contexts, aimed at understanding and improving learning and the environments in which it takes place (Siemens, 2013). This approach leverages data to provide insights that can shape pedagogical choices (Gasevic et al., 2019). For example, analytics might reveal patterns in how students interact with online resources, allowing educators to adjust the course design, or it could identify students who may benefit from early intervention based on engagement metrics.

Educational data mining, on the other hand, focuses on applying machine learning techniques – such as clustering, classification and regression – to uncover patterns and insights from educational data sets (Romero & Ventura, 2020). For instance, analysing data from student discussion forums might reveal links between social dynamics and academic performance, or predictive models might show how participation in certain online activities correlates with retention or achievement.

GenAI brings a significant advancement to both fields. With advanced neural networks, it can model the complexities of human learning within varied environments in ways that traditional techniques struggle to achieve (Sghir et al., 2023). Transparency is essential when using GenAI-driven analytics, allowing both educators and students to understand the basis of decisions made by these tools (Murugesan & Cherukuri, 2023). Explainable artificial intelligence (XAI) provides one approach to this transparency.

XAI comprises techniques that make the reasoning behind AI predictions and decisions more interpretable (Gunning et al., 2019). Tools like Google's Explainable AI offer machine learning models that allow users to see how

variables contribute to predictions, fostering an understanding of the model's decision-making process.

Some ways that you can begin to use GenAI-powered analytics in your own practice include:

- Natural language processing tools can be used to automatically monitor online discussion forums. These tools can analyse the sentiment, tone and semantics of student posts to identify emerging conflict, gaps in understanding or productive dialogue. Based on these insights, you can choose to intervene with targeted guidance or connect students for peer learning.
- Predictive analytics algorithms can be trained on historical student data to build models that identify at-risk learners early in your course. When the model forecasts potential academic struggles for a current student, you are notified to take preventative actions like one-on-one tutoring or additional scaffolds. This enables just-in-time interventions personalized to each learner's needs.
- After group assignments or collaborative tasks, learning analytics dashboards can summarize metrics on participation equality, sentiment and concept coverage to evaluate group cohesion in your course. You can redraw poor performing groups for future projects or use the insights to better design the collaborative experience next time.
- At the end of a course, GenAI can be used to analyse evaluation reports to highlight key trends in participation, engagement and outcomes. These reports can highlight areas of success to repeat as well as suggestions to enhance your future delivery. Sharing such macro-level insights with students also enables self-reflection on their own learning journeys.

GenAI in learning analytics can offer useful insights into student engagement and progress, helping to inform more tailored support. However, it should complement – not replace – human-centred pedagogy. As educators adopt these tools, it is important to prioritize pedagogical intent and ethical use. Without critical oversight, analytics may compromise student privacy, misrepresent learning or expose data to commercial exploitation. Transparent implementation is essential so that these tools serve educational rather than corporate interests. For this integration to genuinely serve education, it must align with the four pillars of this book: Student-Centeredness, Trust, Relevance and Agency. By prioritizing student-centeredness, we enable learners to engage with valued

knowledge, quality, standards and expertise, ensuring analytics work in their best interests rather than merely as monitoring tools. Trust requires that we foster an open environment where students feel safe to take risks and express themselves authentically, unimpeded by concerns about surveillance. Emphasizing relevance means connecting insights from analytics to students' lives and wider social contexts, making the learning process meaningful and applicable. Finally, by focusing on agency, we empower both learners and educators with autonomy and voice in shaping their educational experiences, keeping the use of analytics aligned with personal and collective goals.

Building on these examples, you might be considering how to integrate GenAI tools into your teaching practices. Here are some practical steps to guide this process:

- **Identify specific student needs**. Start by pinpointing the particular challenges or requirements of your students. Do they need more tailored feedback, or are they struggling with specific concepts? Is it essential for them to fully grasp a topic before moving on? Understanding these needs will help you select the most suitable GenAI tool.
- **Research available GenAI tools**. Not all GenAI tools serve the same purpose. Some may focus on specific subjects, like mathematics or language learning, while others offer broader support. Industry news and academic journals, such as EdSurge or the *International Journal of Artificial Intelligence in Education*, feature case studies that can help you identify tools likely to be effective for your teaching context.
- **Pilot before full integration**. Before fully implementing a GenAI tool, consider trialling it in a controlled environment. This could be a short pilot activity, with feedback from students or through your own evaluation.
- **Gather student feedback**. During the pilot phase, collect insights from your students on the tool's effectiveness and areas for improvement. Their perspectives are invaluable for understanding how well the GenAI tool supports their learning experience.
- **Check compatibility with existing systems**. For a smooth integration, look for GenAI tools that align well with your current learning platforms, such as Blackboard or Moodle. A tool that does not function seamlessly with your existing setup is less likely to gain traction.
- **Consider ethical implications**. Introducing new technology that collects student data calls for thoughtful consideration of privacy and ethics.

Verify that the GenAI tool aligns with data privacy standards and refer to Chapter 6 for further discussion on ethical considerations around GenAI in education.

- **Combine human connection with technology**. While GenAI can aid personalization and offer real-time feedback, maintaining the human aspect of teaching is essential. The insights and mentorship you bring to the classroom are irreplaceable, and GenAI should be seen as a supplement rather than a replacement for direct interaction.
- **Share and collaborate**. Connect with other educators who are also using GenAI in their teaching. Exchanging experiences, discussing challenges and sharing solutions can enhance your own implementation and help you make the most of GenAI's potential.

Exercises

Exercise 4.1: exploring pedagogical applications of GenAI

Suggested time: 15 minutes.

Learning outcome: understand the key pedagogical applications of GenAI tools in higher education and their role in personalizing learning experiences.

Description: reflect on a recent course where a GenAI tool might enhance the learning experience by providing tailored resources. Identify one specific area where students might benefit from more personalized support, such as curated reading lists, topic-focused practice questions or instant feedback on assignments. Choose a GenAI tool, ideally one mentioned in this chapter, and document how it could be used within your context. Be specific about where and how you would integrate the tool, any practical challenges that might arise and the potential benefits for students' engagement and learning.

Examples:

- If teaching a literature module, consider using a tool like Research Rabbit to generate reading lists that align with each student's interests, based on their previous selections or research topics.
- For a science module with highly technical concepts, a GenAI feedback tool like Gradescope could provide instant, topic-specific feedback on topics that were causing difficulties, helping students to identify areas for improvement before the next class.

Exercise 4.2: Evaluating GenAI for Inclusivity in Collaborative Learning

Suggested time: 20 minutes.

Learning outcome: critically assess the use of GenAI in supporting inclusive, collaborative learning environments.

Description: think about a group project or collaborative activity in one of your courses. Consider how a GenAI tool might support inclusivity by helping all voices to be represented and by addressing any accessibility needs. Identify three features or functionalities a GenAI tool would need to support these goals and briefly evaluate how well an existing tool could meet these needs. Document both the benefits and limitations of the tool in this context, focusing on how it could support a fair and inclusive environment.

Examples:

- For a discussion-heavy humanities module, a tool like Otter.ai could provide real-time transcription, allowing all students – including those with auditory processing challenges – to follow along and engage fully in group discussions.
- In a collaborative project setting, consider using a GenAI tool like ChatGPT within a shared document platform (e.g. Google Docs). The tool can act as a group brainstorming assistant by generating prompts or summarizing group contributions, helping students build on each other's ideas without a single person dominating. This promotes balanced participation and gives quieter members a way to contribute through suggestions they might otherwise hesitate to share directly.

Exercise 4.3: Identifying practical GenAI tools for self-directed learning

Suggested time: 25 minutes.

Learning outcome: Identify and evaluate GenAI-powered tools for practical use in self-directed and collaborative learning.

Description: identify a course where you believe a GenAI tool could enhance students' independent learning. Choose one GenAI tool from this chapter that might be suitable and document its main features in relation to your specific module needs. Write a brief evaluation (three to four sentences) focusing

on how easy it is to use, the potential benefits for students' learning and any considerations, such as data privacy or over-reliance. Conclude by outlining how you would introduce the tool to students, highlighting its advantages but also addressing its limitations.

Examples:

- For a research-based module, a tool like Elicit can help students identify relevant articles by generating summaries and recommending related sources, which can help them efficiently cover background research.
- In a technical module requiring skill development, a tool like GitHub Copilot could support students in learning programming by providing instant coding assistance and error-checking, helping them troubleshoot and learn through immediate feedback.

Further Reading

This further reading section suggests publications that offer deeper insights and evidence to expand your understanding of the pedagogical applications of GenAI tools discussed in this chapter. Engaging critically with these texts will provide valuable perspectives on integrating GenAI responsibly to support both teaching and learning.

Ouyang et al. (2022) present a comprehensive literature review that examines both the opportunities and challenges of artificial intelligence in higher education. The authors provide a balanced discussion on issues such as privacy, data bias and the changing role of educators, encouraging thoughtful reflection on governance practices when exploring GenAI capabilities.

Kim et al. (2020) explore the rise in online education and the associated increase in demand for AI teaching assistants. Using an online survey, this study investigates students' views on AI teaching assistants in higher education. A central finding is that students' perceived utility of AI teaching assistants and ease of communication with them play critical roles in their acceptance. The study also highlights the need for further research into the nuanced learning experiences shaped by AI assistants.

Williamson et al. (2020) offer a critical examination of the datafication of teaching in higher education, discussing how metrics, data systems and algorithmic methods are increasingly applied to measure, monitor and manage educational processes. The authors argue that datafication may

reinforce competitive logics, surveillance and inequality, addressing concerns such as the commodification of student data and the reduction of pedagogy to quantifiable metrics. Despite these critiques, the paper also considers possibilities for ethical and equitable uses of data, advocating for changes that align data practices with democratic rather than market-driven aims.

Finally, Chan (2023) provides a framework for developing AI policy in higher education, based on stakeholder perceptions. The proposed 'AI Ecological Education Policy Framework' comprises three dimensions – pedagogical, governance and operational – and offers a structured approach to understanding the implications of AI integration. This framework can guide universities in implementing AI technologies in a way that supports both pedagogical goals and responsible practice, underscoring the need for collaboration among all stakeholders to foster a supportive AI environment.

Summary

In this chapter, we have examined the potential of GenAI tools to support and enrich teaching and learning across various domains in higher education. We explored intelligent tutoring systems that offer personalized and adaptive support, collaborative platforms that foster inclusivity, and AI-powered analytics that reveal insights into student engagement and learning behaviours. GenAI presents numerous opportunities to enhance learning experiences in targeted and meaningful ways.

However, integrating these tools responsibly requires careful consideration and continuous oversight. As educators we must assess where GenAI can assist without displacing the relational and human-centred aspects of teaching that are crucial for learning. Ethical considerations, including data privacy, security and transparency, remain central to the effective use of GenAI. The goal is for GenAI to complement rather than overshadow the invaluable human qualities of education – empathy, intuition and personal engagement.

While this chapter focused on the pedagogical applications of GenAI, assessment is another essential area in which AI may have significant implications. In Chapter 5, we turn to GenAI's role in assessment and evaluation, exploring innovations such as AI-driven formative assessments that provide tailored feedback and multimodal evaluations that assess competencies beyond written assignments.

5

Assessment and GenAI Tools

Learning Outcomes

In previous chapters, we explored the potential benefits and limitations that GenAI tools bring to teaching and students' classroom experiences. This chapter examines the impact GenAI tools might have on assessment practices in higher education, considering both opportunities and challenges. While GenAI can raise issues of fairness, transparency and validity that need careful attention, it also offers possibilities for rethinking assessment. For example, could these tools help design individualized assessment tasks or provide more detailed feedback? Alternatively, if we ban GenAI, do we risk losing authenticity in assessment, knowing that these tools will be common in the professional settings our students will enter?

Addressing these questions requires thoughtful reflection on the purpose of assessment and the roles of educators and students. This chapter begins with a review of assessment literacy and then explores the practical implications of GenAI on each phase of the assessment life cycle. As with previous chapters, we assess both the capabilities and limitations of GenAI, considering the value of change as well as the potential challenges.

The chapter includes exercises and examples to help you critically evaluate scenarios for integrating GenAI in assessment. By the end, you should be able to make informed, balanced decisions on the role of GenAI tools in assessment and feedback to support student learning while addressing related challenges.

By the end of this chapter, you will be able to:

1. Describe the influence of GenAI tools on assessment practices in higher education.
2. Design assessment tasks that align with the presence and potential of GenAI tools.

3. Critically evaluate the use of GenAI in grading and providing feedback to enhance student learning.

A Review of Assessment Basics

Assessment design is often a complex subject, even without the addition of new tools that might influence our understanding of academic standards and student performance. For centuries, assessments have relied on traditional methods, such as essays and controlled exams under timed conditions, which remain commonplace across higher education. While these approaches are not inherently flawed, they were developed for a pre-digital era, so it is valuable to reconsider their relevance in the twenty-first century by revisiting the fundamentals of assessment design.

Since Biggs introduced the concept of constructive alignment (Biggs, 2003), most higher education systems now expect assessment design to start with a clear focus on learning outcomes. Learning outcomes define what students should be able to accomplish by the end of a course, establishing both the goals and the standards for the course. Assessment tasks, in turn, should be constructed to allow students to demonstrate achievement of these outcomes, ensuring alignment between learning aims and assessment practices. When this alignment is evident, an assessment is considered valid.

Generally, educators in higher education have substantial flexibility in developing learning outcomes and designing assessment tasks, though outcomes may need to be set well in advance, making rapid changes challenging. For any given set of outcomes, there are multiple task types that can offer valid ways for students to show their learning, as illustrated in Table 5.1.

Assignment tasks also need to be secure, meaning that the teacher must be confident the student has completed the work independently (Dawson, 2020). Security upholds academic standards – our collective expectations regarding the achievement of learning outcomes and the level of attainment, as reflected in grades. Most universities implement policies promoting 'academic honesty', including penalties for students perceived to be circumventing the intended work process. While this seems straightforward, discussions around standards and grading often become complex, as both require interpreting criteria and making subjective judgements (Bloxham et al., 2015; Lloyd & Forsyth, 2024). Implementing measures to maintain security and discourage cheating can also present challenges.

Table 5.1 Examples of Learning Outcomes with Aligned Assessment Tasks for Demonstrating Student Achievement

Sample learning outcome	Indicative assessment tasks
Define and apply foundational concepts in classical mechanics.	• Multiple-choice questions covering a wide range of key concepts.
Analyse complex texts using examples from American literature in the early twentieth century.	• Essay addressing a specific question, using a range of early twentieth-century texts in a traditional academic style. • Debate presenting opposing views on a selected question. • Presentation summarizing key ideas for a targeted audience. • Poster visually conveying core ideas to an intended audience.
Analyse and solve electrical engineering problems in domestic settings.	• Multiple-choice questions on diverse scenarios typical of domestic electrical systems. • Case study analysis of a complex, electrical installation.
Collaborate effectively within a team to solve a human resources management challenge.	• Project report detailing team-based problem-solving strategies. • Group presentation summarizing the team's approach and findings.

A final criterion for effective assessment design is fairness: all students should have an equal opportunity to demonstrate their achievement of outcomes. This includes access to resources and equipment, as well as considerations around the hidden curriculum, expectations and prior knowledge (Campbell, 2022; Forsyth et al., 2022; Nieminen, 2022).

In recent years, the concept of authentic assessment has gained momentum as a meaningful approach to evaluating student achievement (McArthur, 2023). Authentic assessment involves designing tasks that mirror broader societal challenges and encourage students to apply their knowledge in practical, relevant contexts. This approach shifts the focus from rote learning and standardized testing to assessments that are valuable, engaging and directly related to students' future careers and personal growth. For instance, engineering students might be tasked with tackling a real community problem, or environmental science students could be asked to create an art exhibition that reflects the local impact of the climate crisis. These tasks foster deeper engagement and help students develop transferable skills like problem-solving, collaboration and critical thinking, all of which are essential in diverse professional contexts.

With the introduction of GenAI tools, there are new opportunities to make assessment more authentic by creating scenarios or critiques that are hard to replicate through traditional copying or rote responses. The potential for GenAI to generate scenario-based tasks also allows educators to align assessments closely with industry-relevant skills. For instance, a GenAI tool could generate a business scenario that requires management students to analyse and propose solutions, or it could craft a policy-based case study for social sciences students to debate. While these tools provide a valuable resource, they also bring considerations around fairness, validity and the relevance of assessments that we will discuss in this chapter. Embracing these technologies thoughtfully could redefine how we approach assessment, making it more responsive, inclusive and directly connected to students' aspirations and the evolving demands of the workplace.

In this chapter, we will explore how GenAI tools align with these three essential criteria – validity, security and fairness – in ways that support the four pillars of this book. By emphasizing student-centredness, we ensure that assessment methods are valid and tailored to individual needs. Trust is reinforced by secure practices that give confidence in academic standards, while relevance connects assessment tasks to real-world skills, making them meaningful in students' lives and future careers. Finally, agency is fostered when fair assessment design provides students with equitable opportunities to showcase their learning. Through this framework, we aim to approach GenAI's role in assessment thoughtfully, balancing innovation with responsibility.

The Challenges and Opportunities GenAI Tools Present to Assessment

Validity

As previously discussed, the validity of an assessment depends on its alignment with the learning outcomes of the course. GenAI tools do not alter this fundamental relationship, but if they enable students to pass an assessment without genuinely achieving the intended learning outcomes, they introduce a significant challenge to educational objectives. This challenge centres on the meaning of 'you will be able to' within a given learning outcome.

Consider the second example in Table 5.1: 'Analyse complex texts using examples from American literature in the early twentieth century.' This outcome may contain several implicit expectations: that students will read and engage

with a selection of complex texts chosen by the teacher, interpret these texts using widely accepted analytical frameworks, and recognize and compare themes based on these frameworks. Additionally, there may be an expectation for students to demonstrate these skills under controlled conditions, such as an unseen, timed and supervised examination.

With GenAI tools, students could potentially summarize the literature, identify themes and make comparisons without fully engaging with the texts themselves. They might produce an essay, debate contribution, presentation or poster without having read the original material. Many educators might consider such an assessment invalid if students achieve it by bypassing the implicit expectations of the learning outcomes. A discussion of what 'will be able to' means in this subject area is necessary to maintain validity in learning outcomes.

There is also a counterargument regarding the potential benefits these tools offer to students who may struggle to engage with all original texts in depth. Should a literature course require students to read every text in its original form, or could watching a well-regarded film adaptation that captures key themes, or reading a critical analysis by a leading scholar, be acceptable alternatives?

The core issue is that assessment expectations are often socially constructed and tacitly held (Bloxham et al., 2015). For assessments to remain valid in an era of GenAI, these expectations need to be made explicit. This requires more than simply banning GenAI tools during assessments; it calls for a deeper examination of what we aim for students to learn and accomplish through university courses.

Security

Once we have clarified what we want students to demonstrate through assessment, the next challenge is to confirm that students are completing the work independently. If a student falsely claims to have completed an assignment on their own, we consider this cheating, and universities penalize students where cheating can be proven.

Traditionally, measures to prevent cheating rely on controlled conditions, restricting the location, time and resources available to students during assessment. For example, a typical examination format might include an unseen topic in early twentieth-century literature, completed in three hours without notes, texts or GenAI tools. Under such conditions, a student's ability to write an essay would likely demonstrate they have met the learning outcome, with the task considered secure. Similarly, a nursing student who accurately performs prescribing calculations in a simulated clinical setting would give confidence in

their readiness for real-world practice, and thus pass the prescribing module. Security, therefore, is often achieved by limiting resources – still a common response to challenges posed by GenAI.

However, this approach has limitations. Certain learning outcomes require students to work over extended periods, allowing for data collection, analysis and reflection. Here, GenAI tools could potentially substitute student engagement and critical thinking at each stage. One common solution to maintain security is the use of detection tools to check students' work for similarities to other published material. While these tools provide some guidance for educators, they come with challenges, including the effort needed to build a strong case for proving cheating. As discussed in Chapter 3, they also do not yet reliably detect text generated by GenAI, which, by its nature, can vary significantly each time it is used. A skilled student can refine prompts to produce work that bypasses these detection systems.

Consider a different kind of learning outcome, such as a group's ability to solve a complex problem in human resources management. We might limit the project's time frame, conducting it over a single day and try to control access to resources. However, this could be impractical and unrealistic, especially in professional contexts where access to GenAI tools may be standard practice. Denying students access to GenAI in such tasks might reduce the authenticity of the assessment, ultimately making it less valuable for both students and future employers. Perhaps, in this setting, GenAI could aid students by quickly providing insights into complex legislation or prior case studies.

As with validity, we need to examine our assumptions about security. What does it mean to mislead an examiner in an era where GenAI is available, and how do we communicate our expectations to students?

Most students come to university with the aim of learning, often accepting the loss of income and, in some cases, bearing significant living and tuition costs. However, some students may seek only the credentials. If a student does not value learning, lacks understanding of the work or feels incapable of doing it, they may seek shortcuts to obtain the qualification. Such actions may not feel like cheating to them; it may simply feel like efficiency.

Consider what constitutes a shortcut. In pre-internet days, it might have involved finding a book in the library with solved problems from the course or collaborating with classmates on assignments. In these cases, students might still have been learning from the process. Similarly, if teachers reused problems each year with minor adjustments, a tradition might emerge where previous students passed on answers, allowing newer students to follow similar methods.

What if a friend or family member reviews and improves an essay? What if a student, feeling unwell, gets help from a friend to finish it? Or if they cannot start, and a friend suggests a website where they can purchase a 'sample' essay for ideas?

Strategies for obtaining help, taking shortcuts or even crossing into cheating have long existed. In the past, students with resources could access more assistance; now, with information more widely available, GenAI has merely made this challenge more visible.

Fairness

To be fair, assessments must provide all students with an equal opportunity to demonstrate their achievement of the learning outcomes. Historically, assessment methods were designed with a narrower, more homogeneous group of students in mind. Those who succeeded in higher education were typically expected to read quickly and accurately, thrive under the pressure of controlled exams, have a firm grasp of the language of instruction and intuitively understand the implicit expectations of their instructors. These traditional assessments assumed a set of skills and attributes that did not account for the diverse needs of today's student population.

Higher education now aims to be accessible and beneficial to a wider range of students, encompassing diverse backgrounds, linguistic abilities, learning preferences and neurodiverse profiles. Assessment design has evolved to meet these expanded needs while maintaining rigorous standards. This shift has led to more varied assessment types and built-in accessibility features – many of which leverage technology. For example, some assessments incorporate screen readers, spelling and grammar assistance, and even extended time, enabling students with specific needs to participate equitably.

GenAI tools introduce both challenges and possibilities in the pursuit of fairness. On the one hand, their availability raises questions around security and validity, as they might facilitate shortcuts or even bypass certain learning requirements. On the other hand, these tools present opportunities to level the playing field for students facing particular challenges. For instance, students working in a second language or those with neurodivergent processing styles may find that GenAI enhances their ability to interpret complex texts, summarize ideas or express themselves more effectively. This additional support could allow them to focus on demonstrating their understanding and critical thinking, rather than being hindered by language barriers or difficulties with written expression.

The question of fairness in assessments, especially with the availability of GenAI, prompts us to reflect on our underlying expectations. Should assessment focus primarily on a student's ability to perform in conventional, controlled conditions, or should it evolve to emphasize the quality and depth of their understanding and engagement? GenAI might allow for more authentic assessments, where students demonstrate their grasp of a subject through mediums that suit their individual abilities and specific learning needs or preferences For instance, a student with dyslexia might utilize GenAI to organize ideas more coherently, or a student studying in a non-native language might use it to navigate complex texts more effectively.

To make assessment genuinely fair, we must consider whether restricting GenAI tools risks excluding students who need this support to participate on an equal footing. As educators, we must balance the value of traditional assessment integrity with an awareness of these tools' potential to enhance inclusivity. This involves setting clear guidelines around GenAI use, ensuring that its role complements rather than replaces core learning outcomes. Through thoughtful integration, GenAI can become a tool for advancing fairness, supporting a wider range of students in demonstrating their achievements authentically and equitably.

Adapting or Designing Assessment Tasks

As GenAI becomes more widely used in higher education, adapting assessment tasks to reflect the presence of this technology is essential. To maintain validity, security and fairness, you must consider how GenAI can support authentic learning while upholding academic integrity. This section explores ways to adapt assessments when learning outcomes cannot be changed, as well as strategies for designing assessments aligned with updated outcomes that explicitly account for GenAI use. By taking a thoughtful approach to assessment design, you can help students engage with both independent and GenAI-supported skills, preparing them for future professional environments.

Designing Assessment Tasks in a Situation Where You Can Change the Learning Outcomes

In situations where it is possible to change learning outcomes, you can create tasks that explicitly integrate GenAI, allowing students to develop

skills that align with current and emerging professional environments. Updated learning outcomes help students understand how GenAI supports specific skills while reinforcing their need for independent reasoning and judgement. Table 5.2 provides examples of modified learning outcomes and corresponding assessment tasks designed to incorporate GenAI for validity, security and fairness.

Table 5.2 Examples of Updated Learning Outcomes and Corresponding Assessment Tasks with GenAI Integration

Original learning outcome	**Updated learning outcome**	**Strategies for working with GenAI in assessment**
Define and apply basic concepts in classical mechanics.	No change: proficiency in core concepts remains essential.	Controlled multiple-choice questions administered in real-time to ensure thorough understanding and knowledge retention.
Analyse complex texts using examples from American literature in the early twentieth century.	Apply knowledge of early twentieth-century American literature to explore themes in philosophy, science, society and culture.	Portfolio compiling reflections on philosophical and cultural insights, culminating in a short in-class essay under controlled conditions.
Analyse electrical engineering problems in domestic settings.	Analyse domestic electrical engineering problems using both traditional and technological tools.	Multiple-choice questions covering a range of issues, and a case study report that requires students to detail problem-solving approaches using GenAI and other tools, with reflection on the effectiveness of each tool.
Work as part of a team to solve a complex problem in human resources management.	Work collaboratively to develop solutions to complex human resources management problems, using GenAI as an evaluative tool.	Team report documenting solution proposals and explaining how GenAI supported various stages of analysis and decision-making. Presentation: one team member leads, with all members answering examiner questions to demonstrate individual understanding of both the GenAI-assisted and team-based contributions.

Table 5.3 Adapting Assessment Tasks for Unchanged Learning Outcomes with GenAI Considerations

Learning outcome	Assessment task	Strategies for working with GenAI in assessment
Analyse electrical engineering problems in industrial settings.	• Controlled written examination with unfamiliar examples. • Multiple-choice questions covering varied scenarios. • Semester-long case study report.	• Use controlled written examinations to assess technical skills in real-time, minimizing GenAI use. • For take-home reports, require a reflective journal where students document their problem-solving approach, reviewed in class to validate their contributions. • Generate unique case details using GenAI to make reports individualized yet comparable for class discussions.
Prescribe medicines safely within the limits of professional qualifications.	• Practical exam with realistic clinical scenarios. • Multiple-choice questions covering varied situations. • Reflective report on decision-making.	• Integrate GenAI to help students research alternative approaches and protocols, requiring them to document how they validated findings. • Include short reflective reports on how GenAI was used in decision-making. • Ask students to complete regular check-ins during the project to discuss case scenarios and evaluate sources consulted.
Apply statistical analysis to environmental data.	• Data interpretation project. • Lab report with structured problem-solving sections.	• Require students to include GenAI-assisted analysis only in designated sections, while explaining their own interpretation of the data. • Assign an oral component where students discuss their analysis approach, fostering ownership of their work. • Have students compare and critique outputs from different GenAI tools in structured discussions.

Work as part of a team to solve a complex human resources management problem.	• 20-page team project report with 10-minute presentation. • Regular in-class writing and review sessions. • Mini presentations throughout	• Provide access to specific GenAI tools for legal research, helping students identify differences between jurisdictions. • Schedule in-class drafting sessions for teams to produce initial content independently and critique one another's work. • Integrate mini presentations to reinforce progressive task completion and help students articulate how GenAI informed their findings.
Develop an interpretative argument on a historical event.	• Essay with an original thesis. • Short oral defence or discussion.	• Assign GenAI to support initial research and prompt exploration but require students to refine their thesis through their own interpretation. • Use short oral defences to clarify students' arguments and ensure they engage critically with the material. • Encourage reflective notes explaining how GenAI-informed research shaped their understanding, submitted alongside the final essay.

Adapting Assessment Tasks for Unchanged Learning Outcomes

It can take time for learning outcomes to be changed in a course. When it is not yet possible to change learning outcomes, assessment tasks can still be adapted to ensure they remain relevant in a GenAI-permissive environment. In adapting these tasks, focus on integrating activities that require students to engage critically with the content, such as reflective journals, in-class sessions and incremental presentations. Table 5.3 provides examples of ways you can adjust a variety of assessments in this manner.

Generating Assessment Tasks

Incorporating GenAI into assessment design enables you to create flexible, contextually relevant tasks aligned with your learning outcomes. This section provides practical examples for generating essay prompts, case studies and multiple-choice questions, helping you enhance student engagement while upholding academic standards. The examples and strategies below illustrate how GenAI tools can be integrated to support critical thinking and active learning.

Generating Essay-Type Questions

Using GenAI to create a selection of essay topics enables you to present students with diverse and meaningful choices within a given learning outcome. To illustrate this, we asked ChatGPT-4o in October 2024 to generate a variety of essay prompts for the following learning outcome:

> Apply your knowledge of American literature in the early twentieth century to explore philosophical, scientific, social and cultural themes.

ChatGPT generated the following example prompts:

- *How does the theme of isolation in early twentieth-century American literature reflect societal shifts of the period?*
- *In what ways does American modernist literature challenge established norms in science and philosophy?*
- *Analyse how early twentieth-century American writers explore the concept of progress and its impact on individual identity.*

These prompts provide students with the opportunity to engage with course material from different perspectives while staying focused on the core themes. To reinforce students' understanding, consider supplementing the essay task with a brief oral presentation. During this follow-up, other students might answer targeted questions, such as 'How would you connect the theme of isolation to another work from this period?' or 'Can you elaborate on the social critique embedded in your chosen text?' This ensures they have internalized the material and can articulate their insights beyond the written essay.

Creating Case Studies

GenAI can also generate unique, scenario-based case studies tailored to specific group-based learning outcomes. Assigning different case studies to each

student group encourages them to tackle unique challenges, thereby fostering collaboration and critical thinking. For example, we asked ChatGPT-4o in October 2024 to create diverse case studies based on the following learning outcome:

> Work as part of a team to develop solutions for complex problems in human resources management.

ChatGPT generated the following distinct scenarios

1. *A mid-sized company facing high staff turnover needs a comprehensive retention strategy that considers employee engagement and career development.*
2. *A large multinational facing legal challenges in workforce diversity and inclusion requires a policy update to ensure compliance across jurisdictions.*
3. *A small tech startup experiencing rapid growth needs to implement an effective onboarding programme to maintain organizational culture.*

Each group could be assigned one of these scenarios to develop a targeted solution, encouraging them to approach specific HR challenges with practical strategies. For assessment, each group would submit a project report that analyses their assigned case in detail, proposing a comprehensive solution and outlining the tools and processes used throughout the problem-solving process. Following the report submission, one team member would present the case findings, and each group member would participate in a Q&A session. During this session, the examiner would select a question for each student to answer, ensuring that all group members demonstrate individual understanding and engagement with the project.

For example, a group working on the diversity policy case study might develop policy recommendations that align with regional legal standards, using GenAI-enabled research to support their analysis. This structure allows students to grapple with authentic HR issues, applying both technological insights and collaborative skills to their solutions.

Generating Multiple-Choice Questions

Multiple-choice questions are a practical way to assess foundational knowledge, especially for learning outcomes requiring students to define and apply core concepts. GenAI can generate MCQs tailored to your course content, offering varied and contextually relevant questions. For this purpose, we used ChatGPT-

4o in October 2024 to generate MCQs aligned with the following learning outcome:

Define and apply basic concepts in classical mechanics.

ChatGPT provided a set of MCQs that covered a range of difficulty levels, including:

1. Basic concept question

What is Newton's first law of motion?

a) *An object in motion will eventually stop if no force is applied.*
b) *An object at rest will stay at rest, and an object in motion will stay in motion unless acted upon by an external force.*
c) *The force exerted on an object is equal to its mass times its acceleration.*
d) *For every action, there is an equal and opposite reaction.*

Correct Answer: *b) An object at rest will stay at rest, and an object in motion will stay in motion unless acted upon by an external force.*

2. Application-based question

A car accelerates uniformly from rest at 3 m/s^2. How fast will it be moving after 5 seconds?

a) *3 m/s*
b) *8 m/s*
c) *15 m/s*
d) *30 m/s*

Correct Answer: *c) 15 m/s*

3. Calculation question

If a 5 kg object is subjected to a force of 20 N, what is its acceleration?

a) *4 m/s^2*
b) *5 m/s^2*
c) *10 m/s^2*
d) *15 m/s^2*

Correct Answer: *a) 4 m/s^2*

4. Concept application question

In which scenario is Newton's third law of motion demonstrated?

a) *A book resting on a table.*
b) *A car moving at a constant speed.*
c) *A rocket propelling upwards as it expels exhaust gases downward.*
d) *An object falling freely under gravity.*

Correct Answer: *c) A rocket propelling upwards as it expels exhaust gases downward.*

5. Theory-based question

What does the term 'inertia' refer to in classical mechanics?

a) *The resistance of an object to a change in its state of motion.*
b) *The force required to stop a moving object.*
c) *The speed at which an object falls freely.*
d) *The time it takes for an object to reach equilibrium.*

Correct Answer: *a) The resistance of an object to a change in its state of motion.*

The output presents a range of types of multiple-choice question, which may help you to generate ideas for more questions, or may just save you time in designing formative or summative quizzes.

Marking

Effective marking is foundational to fair and constructive assessment. When done well, marking allows students to see clearly how their work aligns with the learning outcomes, where they have met or exceeded expectations, and where they have opportunities to improve. Effective marking should be transparent, with clear criteria that communicate to students exactly what is being assessed and why. It should also be consistent, ensuring that all students are judged fairly across the same standards, and reflective of the intended learning outcomes of the module. This helps maintain the integrity of grading, builds trust and enhances students' engagement with feedback by providing them with specific insights into their strengths and areas for growth.

GenAI tools present a potential to streamline this process by assisting in creating well-defined, consistent marking criteria. With the right prompts, GenAI can generate detailed rubrics that are aligned with specific learning outcomes, ensuring that expectations are communicated explicitly to both students and assessors. GenAI can also help mitigate subjective variation in marking by standardizing criteria, which can be especially useful in assessments that require complex judgement, such as long-form assignments or collaborative projects. However, to be genuinely effective, GenAI-generated rubrics must be carefully tailored to the specific subject matter, teaching goals and contextual

requirements of the course. A significant challenge lies in ensuring that these rubrics do not lose the nuance that comes from experienced educator judgement; criteria generated by GenAI, while consistent, may not fully capture the depth of human insight that is essential for evaluating higher-order critical thinking or creativity.

In objective assessments – such as multiple-choice questions and mathematical problems – marking criteria are straightforward, as the correct answers are typically embedded within the questions themselves. For example, when generating multiple-choice questions (as discussed above), GenAI can provide a correct answer key alongside the questions. This built-in answer guide allows for a streamlined grading process, ensuring accuracy in scoring objective assessments without additional rubric requirements.

For long-form assignments such as essays and reports, where more complex evaluation is required, we trialled both holistic and analytical grading approaches. The holistic approach assesses the assignment as a cohesive whole, gauging how well the overall submission meets the learning outcomes, while the analytical approach examines specific components within the work. To develop a holistic rubric for an undergraduate essay in an American Literature course, we used the following prompt using ChatGPT-4o in October 2024, the results of which are shown in Table 5.4.

> Create a comprehensive grading rubric for a university-level essay in an undergraduate 'American Literature' course. The rubric should cover these areas: clarity of thesis and argument, quality of evidence and analysis, structure and organisation, writing style, critical thinking and adherence to academic conventions. Use five grading bands (Fail, Pass, Good, Very Good, Excellent) and provide specific descriptors for each.

In this instance ChatGPT-4o provided a comprehensive starting point, especially for identifying levels of achievement in structure, evidence and clarity. However, its language needed refinement to fit specific course expectations. For instance, terms such as 'insightful and nuanced thesis' could be rephrased to specify what 'nuanced' means in this context. Additionally, descriptors like 'basic' and 'generally clear' could be elaborated to clarify the performance gap between 'Pass' and 'Good'. Also, while the tool includes technical criteria, it could better capture evaluative language that differentiates critical thinking from mere summarization.

For assessing team-based projects, we tested ChatGPT-4o's ability (in October 2024) to create criteria that balance individual contributions with overall team performance. The assignment requires students to address a complex problem

Table 5.4 Sample Holistic Rubric Generated by ChatGPT-4o for an American Literature Essay, Detailing Criteria at Each Grade Level

Criteria	*Fail*	*Pass*	*Good*	*Very Good*	*Excellent*
Thesis and Argument	*No clear thesis: argument lacks coherence.*	*Basic thesis; limited argument flow.*	*Clear thesis with logical argument.*	*Strong, well-developed thesis and argument.*	*Insightful and nuanced thesis; cohesive and compelling argument.*
Evidence and Analysis	*Minimal or irrelevant evidence.*	*Some relevant evidence, limited analysis.*	*Relevant evidence with adequate analysis.*	*Strong, well-integrated evidence and analysis.*	*Exceptional, well-integrated evidence; analysis demonstrates critical insight.*
Structure and Organization	*Disorganized; ideas lack flow.*	*Basic structure, lacks coherence.*	*Logical structure with clear progression.*	*Cohesive and well-organized, flows smoothly.*	*Exceptionally organized; seamless flow and progression.*
Writing Style and Clarity	*Poorly written, numerous errors.*	*Basic clarity, some errors.*	*Clear and readable, few errors.*	*Polished and articulate, minimal errors.*	*Engaging, professional quality; free of errors.*
Critical Thinking	*Lacks original thought.*	*Limited critical thinking.*	*Shows some original thought.*	*Demonstrates strong critical thinking.*	*Highly original, reflective of advanced critical thought.*
Formatting and Referencing	*Fails to meet academic conventions.*	*Meets basic formatting requirements.*	*Generally correct formatting and referencing.*	*Accurate, follows all conventions.*	*Flawlessly formatted and referenced.*

Table 5.5 Analytical Grading Rubric Generated by ChatGPT-4o for a Human Resources Management Project, Detailing Expectations at Each Level

Criteria	*Fail*	*Pass*	*Good*	*Very Good*	*Excellent*
Teamwork and Collaboration	*Poor collaboration, lack of communication*	*Basic teamwork, limited collaboration*	*Good collaboration, clear communication*	*Strong collaboration, effective communication and conflict resolution*	*Exceptional teamwork, highly effective communication, innovative problem-solving*
Understanding of HR Concepts	*Lacks understanding, major gaps*	*Basic understanding, some gaps*	*Solid understanding, minor gaps*	*Clear and thorough understanding*	*In-depth and sophisticated understanding, insightful application*
Quality of Proposed Solutions	*No viable solution, lacks relevance*	*Basic solution, limited feasibility*	*Practical solution with relevant details*	*Well-researched, effective solution*	*Innovative, thoroughly researched solution with strong practical implications*
Report Presentation and Clarity	*Disorganized, difficult to follow*	*Adequate structure, lacks clarity*	*Well-organized, clear presentation*	*Very clear, well-structured, minimal errors*	*Exceptionally clear, professionally structured, error-free*
Individual Contribution	*Little to no individual input*	*Basic individual input*	*Clear individual input, relevant contributions*	*Strong, relevant individual input*	*Outstanding individual contribution, clearly enhances team's work*

in human resources, presenting their solution in a project report and group presentation. The prompt used was:

> Create a marking rubric for a team-based project on human resources management for second-year undergraduates. Criteria should include teamwork, understanding of HR concepts, quality of proposed solutions and report presentation. Grades should range from Fail to Excellent, with descriptions of expectations for each level.

As can be seen from the output in Table 5.5, ChatGPT-4 provided a solid rubric, especially for collaborative elements and problem-solving criteria. However, the rubric could be improved with criteria capturing the impact of individual contributions on group dynamics, as well as broader societal relevance. For example, 'Good' and 'Very Good' levels in the 'Quality of Proposed Solutions' could specify what constitutes a 'practical' versus a 'well-researched' solution. Furthermore, descriptors for teamwork could include phrases such as 'effective delegation' or 'consensus-building', which offer clearer insight into team dynamics. Including feedback notes, such as comments on each member's unique input, could further support transparency in group assessments.

As these examples illustrate, ChatGPT-4 demonstrates promising potential for helping educators create initial grading rubrics across various assessment types. The structured outputs generated by GenAI can streamline the marking process, providing a reliable foundation that addresses key assessment criteria while saving time. However, despite its efficiency, GenAI's rubric outputs still need careful refinement to align fully with the unique aims and pedagogical nuances of your course. The examples show how critical it is for educators to apply their expertise when tailoring these rubrics – ensuring they encompass not only technical accuracy but also the analytical depth, collaborative skills and originality that are integral to meaningful student assessment.

Feedback

Effective feedback is an essential part of the learning process, guiding students in both understanding their current performance and also in identifying ways to improve and apply what they have learned in future contexts. To be genuinely impactful, feedback must be specific, actionable and tailored to the student's level and stage of study. Ideally, feedback will address the core strengths of a submission while identifying concrete steps the student can take to deepen or

refine their knowledge and skills. Furthermore, feedback in later stages of a qualification should support students' professional development, bridging the gap between academic achievement and career readiness.

GenAI tools offer promising opportunities to support this by generating structured, diverse feedback that can save time and enhance consistency across assignments. These tools can generate comment banks aligned with specific grading rubrics or assessment criteria, making feedback more streamlined. However, using GenAI for feedback also raises questions about accuracy, personalization and security. When deploying GenAI to generate feedback, you will need to closely monitor its output to ensure that it is clear, relevant and practically useful for students.

To illustrate, we asked ChatGPT-4o (October 2024) to generate a feedback bank based on the rubric we created for a team-based human resources management problem-solving report (Table 5.5). We instructed ChatGPT to generate feedback comments that were non-judgemental in tone, focusing on constructive suggestions rather than grade justification. The prompt we used was:

> Using the rubric provided, create a bank of feedback comments suitable for students at the end of their qualification, with a focus on suggestions for future development and professional application.

Table 5.6 presents a sample of the generated feedback, which we edited before adding here to ensure it meets the practical, specific requirements of high-quality feedback.

This generated feedback provides a broad overview of areas for improvement and professional development but could benefit from added specificity to increase clarity and actionability. For example, feedback around 'improving critical thinking skills' would be more helpful if it included direct strategies, such as 'practice evaluating multiple perspectives on a problem to strengthen analytical depth'. Moreover, certain comments could be tailored further to the specific context of the assessment to ensure they address the particular needs and challenges encountered by each student. Incorporating these refinements would make the feedback more practical and supportive for students' academic and professional growth.

To add this level of detail, you might first mark the work, listing targeted bullet points on the student's strengths and areas for improvement. With these points in hand, you could then prompt a GenAI tool, to translate them into polished feedback. This process allows you to define the core aspects of feedback while

Table 5.6 Human-Edited Generated Feedback Comments for Team-Based Human Resources Management Problem-Solving Report

Assessment Criterion	***Generated Feedback Comments***
Collaboration and Teamwork	*'Consider working on enhancing your communication skills for better teamwork. In a professional setting, clear and effective communication is key to successful collaboration.'* *'In future projects, try to take on more leadership roles within your team. Showing initiative and guiding your team can be a valuable asset in your career.'* *'You have demonstrated good teamwork skills. Going forward, continue to build on this strength by actively seeking diverse perspectives within your team.'*
Understanding of the Problem	*'To deepen your understanding of human resources management issues, engage more with current industry trends and case studies. This will prepare you for real-world HR challenges.'* *'Your grasp of the problem is commendable. Continue to hone this analytical skill, as it will be invaluable in diagnosing and solving workplace issues.'*
Quality of Analysis and Solution	*'Focus on developing more critical thinking skills. In your career, you'll often need to analyse complex situations and propose effective solutions.'* *'Your ability to analyse and solve problems is on the right track. Further develop this skill by challenging yourself with more complex scenarios in your field.'*
Report Presentation and Writing Quality	*'Improving your report writing skills, including attention to detail and clarity, will significantly benefit your professional documentation and communication.'* *'Your writing skills are impressive. Maintain this level of professionalism in your reports as clear communication is crucial in any business environment.'*
Use of Evidence and References	*'In your future role, always back up your proposals and reports with strong evidence and research. This adds credibility to your work.'* *'Your effective use of evidence and references is a strength. Continue to apply this rigor in your professional life, as it will support your arguments and proposals.'*
General Professional Development	*'Consider engaging in continuous learning opportunities, such as workshops or certifications, to stay updated in your field.'* *'Networking within your industry can provide valuable insights and opportunities. Make connections and learn from peers and mentors.'* *'Reflect on your experiences in team projects and consider how these lessons can be applied in a workplace setting to enhance team dynamics and project outcomes.'*

using GenAI to enhance clarity and provide actionable suggestions. Reviewing the tool's output is still essential to ensure it aligns with your intended tone and accurately supports each student's academic and professional growth.

It is always wise to check your institution's policy on GenAI use and ensure your practices align with these guidelines. Keeping transparent communication with students about how GenAI is being integrated into feedback and assessment is equally crucial. In doing so, you can foster trust and clarity around the role of AI tools in their learning journey. While GenAI offers substantial support for creating feedback banks or refining bullet points into coherent comments, we strongly recommend avoiding the automatic grading of long-form assessments without thorough quality assurance from you, the educator. Your expertise is essential in ensuring that all feedback remains accurate, nuanced and genuinely supportive of each student's academic growth.

Exercises

Exercise 5.1: Reviewing an Existing Assessment Task Using GenAI

Suggested time: 15–20 minutes.

Learning outcome: identify the influence of GenAI tools on assessment practices in higher education.

Description: select an assessment task that you have previously designed or graded and use a GenAI tool to try to answer it and find the strengths and limitations of a GenAI tool. This exercise aims to help you evaluate how GenAI might impact students' approach to the task and consider any adjustments needed to align with modern standards and ensure fairness. Experiment with various prompts to simulate the types of responses GenAI might generate for students and reflect on any implications this has for the assessment's integrity.

Examples:

- If reviewing a literature analysis assignment, use a GenAI tool to generate a sample analysis based on the task prompts. Assess whether the response meets expected critical depth, if key themes are correctly analysed and if the response relies on sources you want students to access.

1. For a science lab report, prompt the tool with something like 'Generate an introductory lab report on enzyme kinetics.' Review if the structure and

terminology align with what you would expect from students at the same level and adjust as needed to maintain authenticity in student submissions.

Exercise 5.2: Designing a New Assessment Task with GenAI Assistance

Suggested time: 20–25 minutes.

Learning outcome: design assessment tasks that align with the presence and potential of GenAI tools.

Description: use a GenAI tool to help create a new assessment task that reflects current technological capabilities and addresses the values of validity, security, and fairness. This exercise provides an opportunity to explore how GenAI can support your assignment design, from suggesting task types to tailoring complexity. The final assessment should engage students in meaningful ways while minimizing potential shortcuts through GenAI use.

Examples:

- In a course on twentieth-century history, prompt the tool to create a case study assessment where students must compare multiple perspectives on a historical event. Review the output and refine the task by specifying aspects of analysis you value, such as referencing primary sources or critically evaluating historiography.
- For a social science module, use GenAI to draft a group project brief that requires application of theoretical concepts to a broader social issue. Modify the task to clarify roles for each group member, ensuring it encourages independent thought and collaboration rather than reliance on GenAI for all content.

Exercise 5.3: Generating Multiple-Choice Questions with GenAI

Suggested time: 15 minutes.

Learning outcome: critically evaluate the use of GenAI in grading and providing feedback to enhance student learning.

Description: experiment with using a GenAI tool to generate multiple-choice questions (MCQs) aligned with specific learning outcomes for your course. This activity will help you assess the quality, level and practical use of

GenAI-generated MCQs in objective assessments, as well as any adjustments needed to fit them into your assessment strategy.

Examples:

- In a first-year biology course on genetics, ask the GenAI tool to generate ten MCQs that test foundational concepts, including distractors that reflect common misconceptions. Review the questions to ensure they are aligned with the course level, adjusting as needed to avoid overly complex language or inaccurate information.

1. For an introductory economics course, prompt the tool with 'Generate a bank of twenty MCQs for basic macroeconomic principles'. Assess the responses for variety and relevance, ensuring that each question offers a unique angle on the topic to foster comprehensive understanding across key principles.

Further Reading

For those interested in building foundational knowledge in assessment design, Forsyth (2022) offers a clear, accessible introduction to principles of assessment in higher education. This resource is particularly helpful for educators seeking practical guidance, as it covers core concepts and practical examples without requiring extensive prior experience in assessment theory.

If you are exploring the broader challenges of inclusivity and fairness in assessment, we recommend the work of Nieminen (2022) and Nieminen and Lahdenperä (2021). These authors provide valuable insights into the barriers faced by disabled students in traditional assessment settings, and they offer frameworks for addressing these challenges to create more equitable assessments. Their research highlights the ways in which institutional practices can inadvertently disadvantage certain groups, advocating for structural changes to promote fairness in assessment design.

For a comprehensive review of inclusive assessment literature, Tai et al. (2021) critically analyse contemporary studies and argue that, despite increasing attention to diversity, higher education assessment practices still lack inclusivity at many levels. This work is particularly useful for educators and administrators looking to understand the existing gaps in inclusive assessment practices and to identify actionable areas for improvement and how GenAI might be used to achieve this.

Finally, for a critical analysis of current GenAI policies in higher education, consider Luo (2024), who examines university policies from a unique perspective on originality in student work. By reviewing policies from different universities, Luo highlights that many institutions view GenAI as separate from students' own contributions, thus positioning it as a potential threat rather than a resource for learning. This analysis invites a reconsideration of originality in the digital age, suggesting that policy could evolve to support a more balanced approach, where GenAI tools are used to enhance learning rather than being restricted solely out of concerns for misconduct.

Summary

In this chapter, we have explored the multifaceted role of GenAI tools across the assessment life cycle – from the initial design of assessment tasks, through grading and feedback, to broader analysis of student performance and outcomes. Educators may use GenAI tools to draft assignment prompts, establish grading criteria and develop feedback, while students might leverage these tools to assist with their submissions. The implications for grading and the evaluation of students' work are considerable, as GenAI could also aid in scoring submissions and generating feedback. However, with these expanded capabilities come essential questions regarding the desirability and ethical dimensions of these practices. Deciding where human input remains indispensable is critical for maintaining meaningful and secure assessments.

The chapter's discussion is underpinned by the four guiding pillars of this book: student-centredness, trust, relevance and agency. First, GenAI's potential to automate parts of the assessment process could free up educators' time to support students more directly – whether by discussing assignments in depth, reviewing drafts or fostering in-class practice on specific skills. This student-centred approach prioritizes direct engagement, aiming to enhance student development where it matters most.

Trust in GenAI tools is a vital consideration. Educators need to assess the reliability and security of these technologies, especially in sensitive areas such as grading and academic integrity. This includes considering the extent to which both students and educators can be expected to use GenAI responsibly, as well as implementing suitable controls to support this gradual integration at a comfortable pace.

The relevance of GenAI tools in higher education is grounded in their increasing presence in the workplace. By exploring these tools in assessment settings, students gain practical experience with technologies they are likely to encounter professionally. Through intentional and guided use in higher education, students can learn how to operate these tools whilst also evaluating their ethical and practical implications.

Finally, agency is a key element in adapting assessments with GenAI. Could students be given the opportunity to select or customize assessment tasks generated by these tools? Could educators devise individualized prompts or choices that reflect the interests and strengths of each learner? Enabling students to exercise judgement about when and how they use GenAI – and encouraging transparency about its use – cultivates a sense of responsibility and self-management.

As GenAI tools become more integrated into educational practices, educators, students and institutions alike must engage in ongoing discussions about how these technologies should be used. Small, deliberate steps with ample reflection will enable thoughtful adaptation, ensuring that GenAI supports rather than undermines academic integrity and learning goals. In the next chapter, we shift our focus from assessment to student experience, examining how GenAI tools can enhance student engagement, personalization and overall educational satisfaction.

6

Enhancing Student Experience

Learning Outcomes

In this chapter, we shift from the pedagogical applications of GenAI tools discussed in Chapter 4 to a broader exploration of how these tools can enhance the overall student experience. Here, we look beyond the classroom to consider how GenAI can support student well-being, engagement and success. From chatbots fostering connection to recommendation systems offering tailored academic and career advice, this chapter examines the various ways GenAI can enrich student life. We will also explore case studies and examples of GenAI's role in areas like mental health, academic guidance and career support, always with a focus on balancing benefits and ethical responsibilities, ensuring that human judgement remains central.

By the end of this chapter, you will be able to:

1. Identify how GenAI tools can support student engagement, motivation and a sense of belonging in higher education.
2. Understand the role of GenAI in providing personalized support across mental health, academic advising and career guidance.
3. Critically assess the ethical considerations, benefits and challenges of integrating

Student Engagement and Retention

Within the dynamic and ever-evolving landscape of higher education, student engagement is more than a mere metric; it is the backbone of academic success and personal development. While access to vast repositories of information has become commonplace in the digital age, the essence of higher education lies in

ensuring students can engage with valued forms of knowledge and apply them to different situations and contexts. It is crucial that educational systems are not just repositories but gateways that uphold the highest standards of knowledge and expertise (Biesta, 2015).

GenAI tools, with their nuanced and adaptive capabilities, have the potential to affect this engagement. As discussed in Chapter 4, they can be designed so that they do not merely present information but also create a tailored learning journey. By understanding individual learning needs and preferences, these tools might be able to mould content to resonate with each student, ensuring that the academic experience aligns with personal interests, thus fostering a deeper intrinsic motivation amongst students (Vaughn, 2020).

An engaging academic experience is not just about content; it is about context. The environment in which students learn is paramount. The rise of GenAI-powered platforms and virtual learning environments has shown promise in creating spaces that champion open dialogue and risk-taking. Here, students can potentially explore, challenge and articulate their thoughts, fostering genuine curiosity and innovative thinking (Selwyn, 2019a).

One of the most compelling strengths of GenAI lies in its ability to bridge the academic with the personal and to facilitate authentic learning experiences (Crawford et al., 2023). Through continuous analysis of students' interactions, these tools can introduce real-world or realistic applications, case studies and scenarios, bringing theoretical concepts to life. This not only enhances understanding but also provides opportunities for students to demonstrate their knowledge in authentic settings (see Chapter 2).

Effective student engagement extends beyond bridging theory and practice or embedding assessments in real-world contexts. Authentic learning experiences – those that resonate personally and hold individual relevance – can deepen understanding and encourage students to invest more fully in their educational journey. This investment is essential for engagement and retention: when students connect meaningfully with their learning, they are more likely to persist (Pedler et al., 2022). GenAI tools, which can support these personal connections, offer promising potential to sustain student motivation and foster a stronger sense of belonging. This sense of connection is particularly valuable, as it can counteract feelings of isolation or disengagement that often lead to academic withdrawal (Gravett & Winstone, 2022).

As discussed in Chapter 4, the adaptability of GenAI tools offers another advantage: early identification of potential challenges. By monitoring engagement levels and academic performance, these tools can flag areas where

students might be struggling. Such early alerts allow for timely interventions, ensuring students receive the support they need, be it additional resources or mentorship (Jokhan et al., 2019).

However, while the capabilities of GenAI are commendable, it is vital to recognize its limitations. The true essence of engagement often lies in human interactions, be it impassioned debates, guidance from an educator or the camaraderie among peers. AI can augment, but it cannot replicate these authentic human experiences (Zhai, 2022).

As we have stressed throughout this book, the role of GenAI in higher education should be one of complementarity. These tools, while powerful, should enhance, not overshadow, traditional pedagogical practices. It is upon educators and institutions to strike this balance, ensuring that the use of GenAI upholds principles of trust, relevance and genuine engagement.

For readers aiming to use GenAI tools to enhance student retention, you might want to explore the use of Knewton, an adaptive learning platform that personalizes digital courses so students can learn more efficiently (Alam & Mohanty, 2022). Knewton's strength lies in its ability to adapt content to individual student needs. By understanding a student's progress and areas of struggle, Knewton can personalize learning paths, making content more accessible and engaging. Such personalization can reduce feelings of frustration or alienation, often linked to academic disengagement. In doing so, consider the following four steps:

1. **Pilot and reflect with Knewton**. Before a full-scale implementation, introduce Knewton in a single course or module. Monitor its impact on student retention rates and gather student feedback. This trial phase will provide insights into how Knewton can be optimally used in broader contexts and may highlight areas requiring adjustment.
2. **Monitor progress and intervene**. Continuously review the analytics provided by Knewton. It offers detailed insights into student engagement and performance. If a student is disengaging or underperforming, timely intervention, backed by Knewton's data, can provide the necessary support and resources to keep the student on track.
3. **Foster collaborative learning**. While Knewton provides tailored learning experiences, the importance of peer interaction remains. Use Knewton's platform to set up group activities or collaborative projects. Building a sense of community and collaboration within the Knewton environment can further support student retention.

4. **Transparency in data usage**. Ensure students are informed about how Knewton uses their data to personalize learning experiences. This transparency fosters trust, ensuring students feel their personal information is being used responsibly and ethically to enhance their learning, a factor crucial for retention.

Student Belonging

At the heart of a fulfilling higher education experience lies the concept of 'student belonging'. This term encapsulates the sense of connectedness, acceptance and integration students feel within their academic community. It transcends mere physical presence or enrolment, crossing into the emotional and psychological alignment students have with their institution, peers and the broader educational journey (Thomas, 2012).

But can we truly quantify the weight of 'belonging'? Numerous studies (e.g. Illingworth & Radhakrishnan, 2023) have highlighted its direct correlation with academic success, retention and overall well-being. When students feel they belong, they are more likely to be engaged, motivated and committed to their academic pursuits (Strayhorn, 2018).

Despite its importance, fostering a genuine sense of belonging can be challenging. Diverse student bodies, differing backgrounds, varying levels of preparedness and the sheer scale of some educational institutions can create feelings of anonymity and alienation (Doo & Bonk, 2020).

As discussed in Chapter 4, one of the primary ways that GenAI might contribute to belonging is by tailoring educational content to individual student needs. Consider an adaptive learning platform that adjusts course materials based on a student's prior knowledge and engagement with the materials, and pace. These tools can also provide students with real-time feedback on their performance, answering queries and offering resources instantaneously in a way that is simply not possible for human educators; this is not a question of human capability but rather the fact that educators might be away from emails, on leave or even sleeping when some of their learners have specific enquiries. Such personalization can contribute to meeting students' needs for 'active learning, timely feedback, relevance and challenge' which Bovill et al. (2011, p. 1) described as key characteristics for first-year curricula.

While GenAI tools offer numerous advantages, it is essential to be aware of their potential drawbacks; in particular, an over-reliance on technology can

sometimes exacerbate feelings of isolation amongst students (Ma et al., 2023). As such, GenAI tools should be used in conjunction with traditional pedagogical methods to ensure a balanced and holistic educational experience. Indeed, there is a pressing need to develop teaching methods that adapt to our modern, digital age. These methods should foster critical thinking and creativity while balancing certainty with flexibility. They should be transparent in their approach, but also adaptable to the ever-evolving relationship between humans, digital tools and the wider environment (Jandrić & Ford, 2022). As we integrate more GenAI into education, it is essential to recognize and navigate the blending lines between human intelligence, machine learning and the world around us.

Readers keen on harnessing GenAI tools to foster belonging should consider a phased approach. You can work up to more extensive applications of the tools by starting with pilot activities, gathering feedback and making iterative changes. Engage students in the process, explaining why you think these activities might be beneficial to their learning and ensuring they are active participants in shaping the tools and platforms that aim to foster their sense of belonging. Doing so will help to further strengthen student agency and engagement and with it belonging (Tice et al., 2021).

Mental Health and Well-being Support

The mental health crisis among students in higher education has reached alarming levels in recent years. Studies indicate that the majority of university students report experiencing significant anxiety, while suicide stands as one of the leading causes of death for this demographic (Lipson & Sommerville, 2020; Owusu-Ansah, 2020). This mental health crisis can significantly impair academic performance and overall well-being and has only been exacerbated by the Covid-19 crisis and the repercussions that this has had for students (Batra et al., 2021). As such there is an urgent need for comprehensive, accessible and effective mental health support systems within higher education.

Recent studies have explored the potential applications of chatbots in guidance and counselling services (Agarwal et al., 2022). Initial findings indicate chatbots may aid in reducing stress and helping individuals cope with professional and personal challenges. Progress in artificial intelligence and cognitive behaviour monitoring systems presents opportunities to evaluate the feasibility and benefits of advanced technologies in guidance and counselling contexts. These advancements in GenAI offer new pathways for providing personalized and

scalable student mental health services. GenAI counselling chatbots, trained on massive psychotherapy datasets, can offer 24/7 support and guidance to students in need (Guraya, 2023). These chatbots can thus provide an easily accessible first line of support, helping identify those students that require urgent care. As discussed in Chapter 3, GenAI can also anonymously analyse student patterns to flag at-risk individuals.

As highlighted throughout this book, a key benefit of GenAI is its ability to offer consistent yet personalized support. GenAI tools can adapt counselling strategies based on individual needs and cultural backgrounds, helping to promote inclusivity (Koutsouleris et al., 2022). Such an approach has been shown to reduce mental health disparities by increasing the personal relevance of interventions (Hall et al., 2021). Additionally, GenAI may have the capacity to deliver some types of care conveniently through mobile apps, websites and virtual assistants. This expands access to support services, transcending barriers like location, availability and stigma.

Recent advances in computer science and data analytics have opened new possibilities for mental health research and care through GenAI technologies. These tools may support assessment, diagnosis and treatment, while improving access to services. However, systems built on biased historical data risk reinforcing existing inequalities. Developing AI that identifies and mitigates algorithmic bias is therefore essential (Timmons et al., 2023). GenAI should complement – not replace – human counsellors and therapists (Hanley, 2021). While it may improve accessibility and reduce stigma, there is a serious risk that institutions will prioritize cost-saving over care, replacing vital human services with AI-driven alternatives. This could reduce the quality of support and make human care in times of distress a privilege for the wealthiest students. Poor-quality mental health provision also carries severe consequences – including misdiagnosis, suicide and lost opportunities to meaningfully improve students' lives. To mitigate these risks, institutions must prioritize the development and implementation of robust frameworks that emphasize human oversight and accountability. These frameworks should clearly delineate the role of GenAI in student mental health support, ensuring it acts as a supplement to, rather than a substitute for, professional care. Additionally, training models on diverse, unbiased data is critical to delivering inclusive, ethical and effective care, while maintaining the person-centred approaches that are essential for fostering trust and belonging.

For educators looking to implement GenAI mental health tools, we recommend first piloting the technologies to assess their value and potential risks within a controlled setting. Surveying students directly for feedback on

usability, accessibility and perceived support quality can provide critical insights from their viewpoint. Monitoring early usage rates can also reveal student demand levels and any adoption barriers needing redress. Transparency about how student data will be used is essential, along with opt-out choices and regular algorithmic audits by human overseers to uphold ethical data practices that avoid biases. In recommending any of these services to our students, we must take care not to overstate capabilities or position the tools as a catch-all solution; rather, we should emphasize they are intended as supports still requiring human guidance for responsible implementation.

After deployment, regular performance evaluations of these tools coupled with continuous model updates to enhance inclusivity and suitability are suggested, particularly analysing risk analysis and triaging accuracy to pinpoint areas needing increased oversight. By first piloting, soliciting feedback, monitoring adoption metrics, auditing algorithms and evaluating outcomes, we can implement AI in a measured way to best serve all students.

Academic Advising and Career Planning

Navigating academic and career pathways is often complex for university students, given the broad and continuously changing range of degree programmes and graduate opportunities. As such, it is vital for students to receive sound guidance to help them make well-informed decisions that match their abilities, ambitions and interests. Yet, many higher education institutions grapple with providing such bespoke support due to constrained advising resources (Lee et al., 2019).

Historically, university careers services have found it challenging to offer tailored guidance on a large scale owing to limited staff resources. The Covid-19 pandemic, while presenting its own set of challenges, prompted a significant shift to online service delivery. This move has facilitated better engagement for students who find on campus services less accessible. As institutions transition back to face-to-face teaching, there is a clear opportunity for careers teams to meld traditional and virtual offerings. Herein lies the potential of GenAI tools.

Recent developments in AI and machine learning suggest promising methods for delivering personalized academic advising and career guidance. For instance, intelligent chatbots, using natural language processing, can conduct advisory conversations adjusted to a student's academic history, interests and goals (Ilić et al., 2021). These systems can sift through extensive datasets, covering degree programmes, career options, student records and expressed preferences, to craft individualized recommendations. Furthermore, predictive algorithms can align

students with courses and careers based on past data. One notable aspect of GenAI advising tools is their ability to adapt, refining suggestions in response to a student's changing situation and graduate goals. They can also factor in diverse considerations, such as family commitments and preferred locations.

LinkedIn's Career Advice tool serves as an illustrative example. Although not explicitly tailored for higher education, it can offer valuable insights to students. Using machine learning, it assesses a user's profile, skills and activity to hint at potential career paths and job vacancies. By tapping into the vast professional data on LinkedIn, it offers guidance based on criteria such as job roles, essential skills, salary expectations and company cultures.

GenAI-driven chatbots hold potential in enhancing career services at higher education institutions. They offer round-the-clock accessibility, ensuring students can seek advice when needed. With their ability to present labour market insights in an engaging manner, chatbots can also broaden students' perspectives and encourage introspection, while the anonymity they offer might resonate with more reserved students.

However, we need to recognize the limits of these technologies. While they bring efficiency and breadth, they cannot wholly capture the depth and nuance of human interactions. Therefore, GenAI tools should complement, not supplant, human advisors. These advisors bring invaluable mentorship, contextual understanding and the nurturing of skills like communication and empathy.

The guidance process is deeply personal and consequential for students, and while GenAI tools can offer efficiency and breadth, they also introduce questions about data privacy, fairness in algorithmic recommendations and potential biases. The trust that students place in their institutions extends to the tools they are advised to use. Universities must ensure that these tools not only provide accurate and helpful advice but also respect the privacy and individuality of each student. As we blend human expertise with technological capabilities, the moral imperative remains: to provide guidance that is in the best interest of the student, devoid of biases and steeped in a commitment to their holistic well-being and post-university prospects.

Ethical Considerations and Challenges

As GenAI tools are inherently data-driven, often relying on substantial datasets encompassing various facets of student information, there exists a heightened

vulnerability related to data security and confidentiality. As such universities must have clear guidelines in place for where any data is stored, who has access and the safeguards in place to prevent breaches or misuse. This becomes especially pertinent given the global nature of many educational institutions, where data might traverse international borders and be subjected to diverse regulations.

As discussed in Chapter 3, bias and discrimination in these tools is a topic of increasing discourse in the wider AI community (Belenguer, 2022; Ferrer et al., 2021), and one which holds specific pertinence when considering the student experience. The algorithms that power these tools are shaped by their training data, and any inherent biases in this data can be perpetuated and even amplified by the AI. This can result in skewed feedback, recommendations or interventions that unfairly disadvantage or favour certain groups of students.

Addressing these issues requires sensitive handling and understanding of individual circumstances, something that GenAI, if not properly regulated, might overlook. The implications of a misjudgement in these areas can be profound. For instance, incorrect career advice based on biased data can misdirect a student's entire future trajectory. Similarly, mishandling matters of well-being, based on skewed or insensitive AI responses, can exacerbate feelings of alienation or distress. It is crucial that universities acknowledge the gravity of these matters and ensure that GenAI tools, while offering scalability and efficiency, do not compromise the depth, nuance and empathy required when addressing such personal issues. This underscores the necessity for a strong ethical foundation and rigorous guidelines when deploying GenAI tools in such sensitive areas of the student experience.

In exploring the use of GenAI tools for amplifying the student experience, establishing effective governance and ensuring diligent human oversight becomes indispensable. Periodic reviews and audits of these tools, specifically tailored to gauge their impact on the student experience, are essential. Such reviews should consider both the tangible outcomes, such as academic performance, and the intangible aspects, like student satisfaction and sense of belonging.

Feedback is a cornerstone of the educational process, and this holds true in the realm of GenAI tools as well. Institutions should foster an environment where students and educators alike are encouraged to share their experiences with these tools. Such first-hand accounts can offer insights into the real-world impact of the tools, highlighting areas of success and potential improvement.

For those overseeing the deployment and management of how such tools might be used in enhancing the student experience, comprehensive training is

highly recommended. This training should cover the ethical, pedagogical and practical aspects of AI in education. Furthermore, decision-making processes involving GenAI tools, particularly those that have a direct bearing on the student experience, should be transparent, allowing for scrutiny and fostering trust. Finally, while these tools can offer efficiency and scalability, the human touch remains irreplaceable. Systems that integrate AI-driven insights but maintain human judgement can ensure that the enhanced student experience remains grounded, meaningful and ethically sound.

Case Studies

Student Retention

A 2022 study by Arqawi et al. investigated how AI might be the key to predicting whether university students continue with their studies. With student dropouts being a persistent challenge in higher education, the team decided to harness the power of both machine learning techniques and more advanced deep learning methods to tackle this concern.

The study used a dataset of 1,100 student records from a leading university, covering variables such as grades, attendance, extracurricular participation and socio-economic factors. The researchers tested twenty machine learning algorithms to evaluate their ability to predict whether students would continue their studies. Among these, the NuSVC model performed exceptionally well, achieving an accuracy of 91 per cent while also being highly efficient in terms of training and testing time. Building on this analysis, the team explored more advanced computational methods, including deep learning techniques, which offered even greater predictive power. Notably, one of these deep learning models achieved an accuracy of 93 per cent in predicting student continuation, demonstrating the potential for such approaches to outperform traditional machine learning algorithms.

Drawing from their findings, the researchers offer several insights. They underscored the potential of deep learning, urging educational institutions to consider its predictive capabilities with regard to predicting and tackling student retention. In embracing this technology, this study also highlighted the need for such tools to be regularly updated, as with student behaviours and demographics being fluid, it is essential that these predictive models evolve in tandem to stay relevant and accurate. They also highlighted the need for clear ethical guidelines and governance when it comes to sensitive areas like

predicting student dropouts. Students should always be in the loop about how their data is being used, with their privacy being a top priority.

Such tools may help teachers and advisors to identify potential issues earlier, allowing them to follow up with in-person support. Those looking at the data must always remember that predictions based on previous behaviours may only be dependable as averages for the cohort and may not apply to individual students.

Supporting Student Mental Health and Well-being

In their 2023 study, Chen et al. investigated the potential of chatbots within classroom environments, emphasizing their potential role in advancing student success and well-being. The researchers conducted a comprehensive survey involving 215 undergraduate students to gauge their perceptions and experiences with GenAI chatbots.

The findings revealed that most of the students expressed comfort in engaging with GenAI chatbots, especially when broaching sensitive subjects. Some of the students showed a preference for chatbots when exploring career prospects, with many respondents valuing the unbiased nature of AI in providing career advice. Similarly, a significant number of students found chatbots a safe space for discussing mental health concerns, citing the absence of judgement as a key advantage.

However, Chen et al. also underscored the need for responsibility and transparency. While students' trust in chatbots is evident, the study argues that institutions have a duty to ensure clarity in how chatbots process and use user data, particularly when handling sensitive topics like mental health. They recommend regular audits and training for GenAI systems to maintain fairness and preclude inherent biases. Furthermore, the research emphasizes the importance of synergistic efforts between educational entities and AI developers. Such collaboration ensures that chatbot tools are continuously refined, catering to the multifaceted needs and evolving preferences of students. In the EU, the AI Act 2024 will probably consider such applications to be 'high-risk' and subject to careful oversight.

Career Counselling

In their 2023 study, Guleria and Sood investigate the applications of explainable AI (XAI) and machine learning (ML) within career counselling. Their

methodology involves the use of educational data mining techniques to design and test classifiers, harnessing sensor data and various student attributes.

Their research underscores the transformative potential of XAI and ML in making the career counselling process more efficient and tailored. A pivotal contribution of their work is the introduction of an AI-powered career guidance system. This system, leveraging natural language processing and student data analysis, aims to provide individualized career recommendations.

The findings from Guleria and Sood's study highlight a marked enhancement in student engagement and precision in career path predictions with the integration of AI-driven models. Importantly, the study accentuates the critical need for transparency in these models, with transparency recommended for educators and students alike, given the roles these tools might play in shaping prospective career decisions.

Exercises

Exercise 6.1: Evaluating ChatGPT for Student Engagement

Suggested time: 20 minutes.

Learning outcome: identify how GenAI tools can support student engagement, motivation and a sense of belonging in higher education.

Description: engage with ChatGPT to assess its potential in fostering student connection and motivation through virtual engagement. This will involve using ChatGPT to simulate discussion prompts or study group dialogues, helping you identify ways this tool can reinforce student motivation and community within a course.

Examples:

- For a first-year humanities course, prompt ChatGPT to generate five discussion questions on a given topic (e.g. 'Generate engaging discussion prompts for a module on Renaissance art'). Observe the relevance and depth of the prompts and consider adjustments to enhance their potential in sparking meaningful interactions among students.
- If teaching a large introductory science course, ask ChatGPT to create motivational messages for students at key points (e.g. 'Create a mid-term

motivational message for students struggling with introductory physics concepts'). Reflect on how well these messages might encourage persistence and a sense of support.

Exercise 6.2: Using Wysa for Personalized Mental Health Support

Suggested time: 25 minutes.

Learning outcome: understand the role of GenAI in providing personalized support across mental health, academic advising and career guidance.

Description: explore the mental health support features of Wysa, a popular AI-driven mental health chatbot, to evaluate how it provides personalised support for common student concerns. This will help you assess Wysa's ability to deliver thoughtful and accessible guidance while identifying any limitations in its responses.

Examples:

- Begin by asking Wysa for guidance on managing academic stress (e.g. 'What strategies do you recommend for reducing stress during exam periods?'). Evaluate the appropriateness and variety of responses, considering how students with different stress levels might benefit from this tool.
- For a scenario where a student feels isolated or overwhelmed, interact with Wysa to simulate responses, noting how effectively it validates feelings and offers practical advice. Reflect on potential gaps where additional human counselling might be needed.

Exercise 6.3: Balancing Benefits and Challenges with Gradescope for Automated Feedback

Suggested time: 30 minutes.

Learning outcome: critically assess the ethical considerations, benefits and challenges of integrating GenAI in student-centric contexts.

Description: examine Gradescope's capabilities in providing AI-driven feedback on assignments, assessing both its benefits in timely feedback and potential challenges in maintaining accuracy and student trust. In doing so, consider when human review might be necessary to supplement automated feedback.

Examples:

- With their permission, upload a sample student assignment (e.g. an essay or lab report) to Gradescope and review the AI-generated feedback. Analyse the strengths and weaknesses in feedback quality, especially in areas that may require nuanced responses. Reflect on when human oversight might improve feedback effectiveness.
- For an assessment with multiple-choice or short answer questions, test Gradescope's automated grading feature. Note any errors or inaccuracies in grading and discuss scenarios where human intervention would be crucial to ensure grading fairness and reliability.

Further Reading

For those interested in leveraging data insights to support student engagement and retention, Hussain et al. (2018) provide a compelling example of how AI-powered analysis of sensor data can identify students who may be at risk. Their findings suggest that early interventions, informed by AI-driven engagement metrics, have the potential to boost student performance and persistence. This research offers a practical foundation for integrating similar strategies into teaching practices, helping educators respond proactively to student needs.

Learning analytics, a field that has informed the development of GenAI tools, remains relevant as we consider enhancing student belonging and well-being in higher education. Ihantola et al. (2019) review the practical applications of learning analytics and discuss its role in personalizing learning experiences. Although distinct from GenAI in its technical scope, learning analytics provides a complementary approach, offering insights that help educators align AI-driven engagement initiatives with broader student experience goals.

For institutions exploring AI-driven career advising, Vignesh et al. (2021) offer a useful conceptual framework for developing chatbots that deliver personalized career guidance. By using natural language processing and analysing student data, these chatbots offer targeted recommendations that can be continuously refined. This framework underscores the importance of critically evaluating both the technological and ethical dimensions of such systems to ensure meaningful support for students' career paths.

Ethical considerations are fundamental when introducing GenAI into student support services. Firat's (2023) research highlights the ethical responsibilities

educators face, especially in relation to data privacy and the need for transparency with students. Firat emphasizes the importance of balancing innovation with student agency, suggesting practical steps to responsibly integrate GenAI while respecting student data. Timmons et al. (2023) also emphasize the need for ethical oversight, calling for diverse data sources, regular audits and ethical governance, particularly when applying AI in sensitive areas like mental health. Their analysis reinforces the need for ongoing vigilance to ensure that GenAI tools enhance, rather than compromise, student well-being and agency.

Summary

Chapter 6 explored the influence of GenAI tools on the student experience in higher education, focusing on how these technologies can enhance engagement, motivation and belonging. Through practical case studies and real-world examples, we examined the potential for GenAI tools to personalize academic support, foster well-being and offer guidance tailored to students' needs. These insights aim to help you identify where and how GenAI tools might complement and enrich your own teaching practices, especially in supporting diverse student needs.

We also addressed the complexities associated with these tools, including the ethical and practical considerations necessary to safeguard student trust and agency. Issues like data privacy, potential biases in recommendations and the importance of transparency and human oversight were underscored. These concerns are essential to managing the benefits of GenAI effectively while ensuring that student support remains rooted in personal and ethical principles.

As we move into Chapter 7, which synthesizes the book's core themes, we encourage you to reflect on how the four pillars – student-centredness, trust, relevance and agency – apply to your practice. This final chapter aims to bring together our discussions, offering a framework for thoughtfully integrating GenAI into higher education, equipped with a balanced view of both its opportunities and its challenges.

7

Conclusions

Learning Outcomes

In this final chapter, we bring together the core themes, insights and practical applications discussed throughout the book to reflect on GenAI's evolving role in higher education. We have journeyed through foundational principles, ethical debates and practical exercises that illustrate how GenAI can be integrated meaningfully into teaching and assessment. Here, we draw these ideas together to consider how they might shape the future, particularly in a rapidly advancing educational landscape.

Through engagement with the book's content and exercises, you should have gained valuable skills and perspectives that align with the following overarching learning outcomes:

1. **Integrate theoretical and practical GenAI applications.** By now, you should be able to synthesize GenAI's theoretical underpinnings with practical applications, creating a balanced understanding that connects foundational concepts with the diverse case studies presented. This integration empowers you to apply GenAI in educational contexts thoughtfully and with purpose.
2. **Anticipate and prepare for future GenAI needs in higher education.** You are now positioned to identify the emerging skills, knowledge and strategies essential for harnessing GenAI effectively in the future. This readiness involves staying current and adaptable, ensuring you are prepared for future developments that align with evolving pedagogical goals.
3. **Critically assess ethical and practical implications of GenAI.** Equipped with a nuanced perspective on the ethical, pedagogical and operational dimensions of GenAI, you should now approach its application with

a critical, informed stance. This ensures that any integration of GenAI upholds ethical integrity, prioritizes student-centred learning and avoids the risks associated with unreflective technological adoption.

A Reflective Summary of GenAI in Higher Education

In this final chapter, we return to the central themes and insights that have guided our discussion on GenAI within higher education. Across the previous chapters, we have explored the capabilities, potentials and limitations of GenAI tools in a range of educational contexts, considering both the practical applications and the ethical complexities inherent in their use.

We began by laying the groundwork with a fundamental understanding of GenAI, distinguishing between basic, rule-based systems and advanced machine learning algorithms that adapt over time. Here, we established how these technologies might tailor educational experiences to meet individual student needs, a recurring theme throughout this book. This foundation led us into Chapter 2, where we examined specific tools in use today and those on the horizon, weighing the efficiencies these tools bring against the skills they require, such as digital literacy, and the risks of overdependence.

Our exploration continued in Chapter 3, where we considered the ethical and social dimensions of GenAI in higher education. This chapter discussed the responsibility of educators and institutions to implement GenAI tools in ways that support, rather than undermine, equity in access and learning quality. By developing an ethical framework, we can ensure that GenAI is used to enhance educational inclusivity without reinforcing existing inequalities. Chapter 4 then examined specific pedagogical applications of GenAI, looking closely at how these tools can be used to enhance classroom experiences, support specific learning needs and engage students more actively. This chapter was grounded in a commitment to student-centred learning, showing how educators can thoughtfully incorporate GenAI to enrich curriculum design, extend personalized support and enhance engagement – always with the caveat that these tools should supplement, not replace, critical human interactions in teaching.

Chapter 5 moved on to consider GenAI's influence in assessment and feedback. Here, we examined how these tools might streamline grading, offer timely feedback and even help create innovative, dynamic assessment tasks that

better align with learning outcomes. This chapter also raised questions about the limitations of GenAI in marking complex, subjective tasks, underscoring the essential role of human judgement in delivering fair, nuanced feedback. Across these discussions, we consistently highlighted the necessity of human oversight, arguing that while GenAI can support and enhance various aspects of teaching, learning and assessment, it cannot substitute for the unique insight and sensitivity that educators bring to their students' learning experiences.

Chapter 6 extended this discussion to the broader student experience, exploring how GenAI can support engagement, belonging and more personalized educational paths. Authentic examples illustrated GenAI's ability to help students take ownership of their learning, while also reminding us of the importance of keeping human connections at the core of educational experiences.

Throughout this book, four guiding principles have framed our discussions on GenAI in higher education: student-centredness, trust, relevance and agency. We have advocated for a student-centred approach that keeps students' interests and learning outcomes central, emphasizing trust between students, educators and institutions as essential for a GenAI-enhanced educational environment. We examined relevance, showing how GenAI's adaptive capabilities can create more meaningful learning experiences without diluting the critical depth higher education demands. And finally, we championed agency – both for learners and educators – as foundational to any successful GenAI integration, ensuring control over how these tools shape the curriculum, teaching and assessment.

In reflecting on these insights, GenAI presents real promise for transforming higher education. Yet, its successful and sustainable integration requires an informed, ethical and principled approach – one that balances technological possibilities with the values that underpin meaningful education. As we look ahead to the future applications of GenAI, we must do so critically, applying what we have learned here to shape an educational environment that is equitable and enriching for all.

Preparing for Future Technology Developments

In contemplating the horizon of technological advancements, educators and institutions face the formidable task of not only anticipating future trends in GenAI but also shaping them. The narrative we have woven throughout this book suggests a proactive stance. By engaging with the development community and

participating in technology governance, educators can lend their critical insights to guide GenAI towards pedagogically sound applications. This involvement ensures that future iterations of GenAI tools are designed with an educator's perspective, incorporating features that facilitate teaching and learning rather than dictate them (Kaplan-Rakowski et al., 2023).

Ongoing professional development emerges as a cornerstone in the educator's toolbox for remaining relevant in a shifting technological landscape. Lifelong learning, a concept we have championed since Chapter 2, is not just a lofty ideal but a practical necessity (De la Harpe & Radloff, 2000). As GenAI technologies evolve, so too must the skillsets of those who wield them in the educational arena. This means that as educators we must commit to a continuous cycle of learning, relearning and unlearning – a cycle that is as iterative as the algorithms we seek to harness.

The pace at which GenAI is advancing calls for a flexible approach to professional development. Educators should be afforded opportunities to experiment with emerging technologies, gaining first-hand experience that is both reflective and practical. This direct approach was exemplified in the case studies of Chapter 4, where educators not only learned about GenAI but also applied it in diverse pedagogical contexts. Such immersive experiences are invaluable for demystifying technology and fostering an adaptable mindset (Oprean & Balakrishnan, 2020).

When it comes to policy development, institutions must be as dynamic as the technologies they aim to govern. Policies that were discussed in Chapter 6 should not be static documents but living frameworks, subject to regular review and revision in response to technological advancements. They must embody principles of equity, access and inclusivity, ensuring that GenAI tools enhance rather than encumber the educational experience (Farrelly & Baker, 2023).

Infrastructure plays a critical role in the support and deployment of future GenAI technologies. As we have seen through the discussions in Chapter 5, the integration of GenAI requires robust digital infrastructure that can handle large-scale data processing and deliver personalized learning experiences. Institutions will need to invest in both hardware and software, but also in the human capital – the IT professionals, data scientists and educational technologists who can bridge the gap between potential and practice (Fullan & Langworthy, 2014).

The agility of institutions in adapting their infrastructure to support GenAI will be a significant determinant of their success. This agility was evident in the case studies from Chapter 3, where institutions that embraced change and innovation were able to leverage GenAI tools to enhance student learning outcomes. In these instances, the infrastructure served as the backbone for

innovation, enabling educators to experiment with and refine their pedagogical approaches (Chatterjee & Bhattacharjee, 2020; Vodenko & Lyausheva, 2020).

Collaboration between educational institutions and technology developers will be pivotal in steering the future of GenAI. Such partnerships, as highlighted in Chapter 2, can lead to the co-creation of technologies that are both innovative and pedagogically relevant. This symbiotic relationship ensures that GenAI tools are designed to meet the immediate needs of educators, while remaining adaptable to future demands (Southworth et al., 2023).

The ethical considerations that we have consistently revisited, most notably in Chapter 3, will continue to be a guiding light for future developments. Institutions must advocate for the ethical use of GenAI, ensuring that the technology is employed to serve educational goals and not the other way around. This advocacy requires a nuanced understanding of the ethical landscape and a commitment to uphold the highest standards of integrity and fairness (Nguyen et al., 2023).

Preparing for future technology developments is not a solitary journey. It requires a collective effort, as we have underscored throughout this book. The sharing of effective practices is essential. Such collaborative efforts can help create a common understanding and shared vision for the role of GenAI in education, paving the way for its responsible and effective integration (Ng et al., 2023).

Preparing for the future of GenAI in education is an ongoing process of anticipation, adaptation and ethical consideration. It requires educators to be lifelong learners, institutions to be agile and forward-thinking, and policies to be inclusive and dynamic. As we have seen through the various chapters of this book, the future of GenAI in higher education is not a distant speculation but an unfolding reality that demands our active engagement and thoughtful preparation.

As you seek out or are offered GenAI applications which may be integrated into existing educational software or presented as new products, you will be faced with practical decisions about whether to test and adopt them. To help with asking questions and making these choices, we have summarized in Table 7.1 some of the potential advantages and disadvantages of a range of different potential applications of GenAI in higher education of the kind we have discussed in the previous chapters. In generating these examples, we have drawn on the key themes of student-centredness, trust, relevance and agency. Only you will be able to decide whether a product will benefit student learning in your own teaching context, and we hope you now feel better equipped to take those decisions.

Table 7.1 Advantages and Disadvantages of Different Educational Applications Using AI

Application	Potential benefits in higher education	Potential disadvantages in higher education
Advanced searches and reviews of scholarly material	Could improve the range of scholarly material used to take decisions in education. Could be useful for students starting on their first attempts at research as they can get going without as much help from teachers and librarians. Allows students to select their own areas of interest.	Lack of specialist support might result in the selection of biased, outdated, irrelevant and low-quality selection of scholarly work and other materials. Possible inaccuracies in outputs. Students may not gain the skills they need for future research activity.
Planning course sessions and activities	Can generate a wide range of interesting and stimulating activities which are enjoyable for teachers and students. Can ensure, via clear instructions in the prompt, that accessibility is considered. May save teacher time.	Plans still need to be reviewed for relevance, quality and absence of bias. Dependence on these activities might lead to teacher detachment from teaching.
Production of multimodal materials: convert text and other data to online learning materials, videos and images	Improve the quality of online resources. May save teacher time. Can ensure that accessibility is considered.	Time will still be needed to check quality, and criteria will be needed for this. Respect for the economic and moral IP rights of authors of original materials will need to be ensured.
Examination design: systems generate questions, tasks and instructions based on parameters input by teachers	Can generate a wide range of interesting and challenging requirements which are enjoyable for teachers to grade and for students to complete. Can ensure that accessibility is considered.	Outputs will need careful review for possible inaccuracies in relation to content, level and standards.

Application	Potential benefits in higher education	Potential disadvantages in higher education
Automated grading and feedback systems, which review student work according to preset criteria	May save teacher time.	Possible inaccuracies. Set-up time can be substantial. Impression of false security with the outputs, which may be misleading or wrong. Who is responsible for what the tool says: the teacher, or the company making the software? What if it is wrong? Illusions of objectivity: in reality, subjectiveness will be built into the criteria, as it is when teachers grade without using technology. Teachers need to engage with their students' work, which shows what they have learned. This interaction about purpose is key to trust between teachers and students.
Adaptive personalized learning: students are presented with different content and questions based on their answers to previous tests or choice of previous study material	Students have unlimited access to the materials, unlike with a teacher. Students can go at their own speed. Students can explore particular interests in more depth. We might be able to ensure that accessibility is considered and adapted to individuals.	Teachers may lose control over what students are learning: the applications may move away from the curriculum or contain inaccuracies. Who is responsible for the options and the quality of material presented in the tool? Can the outputs be trusted? It may be difficult to develop applications which encourage deep learning and critical thinking: they may be restricted to knowledge acquisition.

Application	Potential benefits in higher education	Potential disadvantages in higher education
Data analytics: Can monitor student participation and attendance in different activities, highlighting changes in pattern which may indicate problems	Might make support available more quickly. Could save teacher and pedagogical support time.	Potential for violation of fundamental rights. Might in some cases amount to a prohibited use of AI. Risk of predictive tools to provide information not necessary and proportional to the educational task. Potential for indirect detection and sharing of personal information concerning private matters (health, family life, lifestyle, etc.)? Who decides what actions are needed in response to the analysis? Who is responsible for the answers?
Personal tutoring/coaching	Each student has a digital application which mimics the work of a tutor, career counsellor or mentor available 24/7.	Possible inaccuracies. Danger of allocating resources to expensive AI systems, instead of employing properly trained teachers. Impression of false security with the outputs, which may be misleading or wrong. Who is responsible for what the tool says: the teacher, or the company making the software? What if it is wrong? GenAI tools outputs tend to focus on high probability scenarios, when higher education encourages thinking about creative and low probability scenarios. Potential for indirect detection and sharing of personal information concerning private matters (health, family life, lifestyle, etc.). Ultimately impersonal: there is some evidence that students are less likely to respond to feedback from computer-based systems than from human teachers.

Exercises

Exercise 7.1: Synthesizing Key Themes with GenAI

Suggested time: 20 minutes.

Learning outcome: integrate theoretical and practical GenAI applications.

Description: use ChatGPT to synthesize the main themes discussed throughout this book, focusing on key areas like student-centred learning, trust, relevance and agency. Begin by prompting ChatGPT to generate concise summaries of each theme, then organize these summaries visually to create a mind map that links each theme to relevant chapters. Conclude by reflecting on how these themes can inform specific teaching practices.

Examples:

- Prompt ChatGPT with 'Summarize the central themes of this book in relation to student-centred learning, trust, relevance and agency'. Use the summary to build a visual mind map, showing connections between each theme and specific chapters.
- Request a summary of 'ethical considerations of GenAI in education, as covered in this book', then draft a brief outline that integrates theoretical insights with practical examples for your own teaching scenarios.

Exercise 7.2: Planning for Future GenAI Skills

Suggested time: 25 minutes.

Learning outcome: anticipate and prepare for future GenAI needs in higher education.

Description: engage in a dialogue with ChatGPT to identify areas for continuing professional development (CPD) that align with the evolving demands of GenAI in education. Begin by uploading or summarizing your CV, then ask ChatGPT to highlight skills and areas where further CPD would be beneficial for integrating GenAI tools effectively in your work. Use its suggestions to create an actionable plan.

Examples:

- Share a brief summary of your CV with ChatGPT (e.g. 'I am a lecturer in higher education with a focus on curriculum design and digital

assessment'). Request feedback with 'Identify gaps in my skills related to GenAI integration in education and suggest CPD areas to address these.'
- Ask ChatGPT, 'Based on my current experience, outline specific CPD activities to help me incorporate GenAI ethically and effectively in my practice.' Use the response to create a development plan, including recommended courses, certifications or workshops, tailored to your background.

Exercise 7.3: Drafting an Ethical Policy for GenAI in Student Assessments

Suggested time: 20 minutes.

Learning outcome: critically assess ethical and practical implications of GenAI in higher education.

Description: develop a brief ethical policy for using GenAI in student assessments by identifying key considerations and principles. This exercise will involve exploring ChatGPT's input to help shape a policy that aligns with ethical standards, balancing transparency, fairness and data privacy.

Examples:

- Ask ChatGPT to 'List ethical considerations for using GenAI in student assessments, including data privacy, fairness and transparency.' Use its response to draft three to four guiding principles, such as 'protecting student data privacy' and 'ensuring clear communication on GenAI's role in assessment'.
- Prompt ChatGPT with 'Generate a short explanation for students on how GenAI will be used in their assessments.' Use this response to draft a section of the policy that helps build trust by clearly explaining the GenAI process.

Further Reading

Sarı et al. (2024) investigate GenAI's potential for teaching and learning, framing it through a lens of critical pedagogy. Their work explores the tension between using GenAI to automate aspects of education and empowering learners, inviting educators to reflect on the sociotechnical influences that shape GenAI's role in higher education. This perspective on GenAI highlights the importance

of maintaining a balanced approach, one that considers both efficiency and educational agency.

In a related exploration, Holmes et al. (2022) survey experts on GenAI in education, finding a consensus on the ethical complexities introduced by AI integration. Key concerns – such as fairness, transparency and inclusion – are acknowledged as foundational principles that must be explicitly addressed in GenAI applications. The researchers note a gap in training for addressing these ethical concerns, suggesting that higher education institutions take active measures to incorporate ethical considerations as part of GenAI implementation rather than relying solely on positive intentions.

A 2021 study by Buçinca et al. serves as a cautionary note on over-reliance on GenAI, especially when tools are used in decision-making contexts. The study found that interventions prompting deeper engagement with AI explanations helped mitigate blind reliance on AI outputs, including incorrect recommendations. However, these interventions were rated less favourably by users, underscoring a critical trade-off between fostering thoughtful AI use and user satisfaction. This balance between promoting informed engagement and user ease is essential for institutions to consider as they integrate GenAI systems.

Kaplan-Rakowski's (2023) study explores GenAI adoption from the teachers' perspective, revealing a mix of enthusiasm and caution. The research underlines the need for participatory governance models in higher education that involve all stakeholders, ensuring that human values are central to GenAI's development and use.

Finally, for a broad overview of GenAI's societal implications, Lee and Qiufan's *AI 2041: Ten Visions for Our Future* (2022) explores AI's impact on work, culture and communication. Their analysis spans technical, social and economic dimensions, prompting educators and institutions alike to consider how AI's evolution will reshape the future of education and beyond. This work is a valuable resource for those looking to reflect on the broader societal context of GenAI.

What Next?

As we draw this book to its close, we invite you, the reader, to take a moment to consider the journey we have traversed together. The landscape of GenAI in higher education is vast and varied, and while we have explored many of its aspects, there remains much ground to cover.

Your next steps are as important as the words you have read. We encourage you to continue the conversation, both within your professional circles and in the broader educational community. Share your insights, challenge the ideas presented and contribute your unique perspective to the evolving narrative of GenAI in education.

Should you wish to further investigate any of the topics we have discussed, or if you are eager to share how you have applied the concepts of this book in your practice, we welcome you to reach out to us. Your experiences, questions and stories enrich the collective understanding and push the boundaries of what we can achieve in integrating GenAI into higher education.

You can connect with us via the contact details provided at the end of this book. Whether it is to offer feedback, seek advice or propose collaboration, your engagement is what propels this work beyond the confines of its pages.

Lastly, we would like to express our sincerest gratitude for your company on this explorative journey. The field of GenAI in higher education is one of constant change, and it has been our pleasure to navigate some of its complexities with you. May the knowledge and ideas you have gathered here inspire you to forge new paths and create impactful educational experiences.

With warm regards and thanks for your company,
Sam and Rachel.

References

Abulibdeh, A., Zaidan, E. & Abulibdeh, R., 2024. Navigating the confluence of artificial intelligence and education for sustainable development in the era of industry 4.0: Challenges, opportunities, and ethical dimensions. *Journal of Cleaner Production*, 437, Article 140527. Available at: https://doi.org/10.1016/j.jclepro.2023.140527.

AbuShawar, B. & Atwell, E., 2015. ALICE chatbot: Trials and outputs. *Computación y Sistemas*, 19(4), pp.625–32. Available at: https://doi.org/10.13053/CyS-19-4-2326.

Agarwal, S., Linh, N.T.D. & Aponte, G.J.R., 2022. Student perception regarding chatbot for counselling in higher education. In: V.E. Balas, V.K. Solanki & R. Kumar, eds. *Recent Advances in Internet of Things and Machine Learning: Real-World Applications*. Cham: Springer International Publishing, pp.263–81.

Alam, A. & Mohanty, A., 2022. Business models, business strategies, and innovations in EdTech companies: Integration of learning analytics and artificial intelligence in higher education. In: *2022 IEEE 6th Conference on Information and Communication Technology (CICT)*, Gwalior, India, pp.1–6. Available at: https://doi.org/10.1109/CICT56698.2022.9997887.

Alsafari, B., Atwell, E., Walker, A. & Callaghan, M., 2024. Towards effective teaching assistants: From intent-based chatbots to LLM-powered teaching assistants. *Natural Language Processing Journal*, 8, Article 100101. Available at: https://doi.org/10.1016/j.nlp.2024.100101.

Alshaigy, B. & Grande, V., 2024. Forgotten again: Addressing accessibility challenges of generative AI tools for people with disabilities. In: *NordiCHI '24 Adjunct: Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction*, pp.1–6. Available at: https://doi.org/10.1145/3677045.3685493.

Anderson, J.R., Corbett, A.T., Koedinger, K.R. & Pelletier, R., 1995. Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4(2), pp.167–207. Available at: https://doi.org/10.1207/s15327809jls0402_2.

Arao, B. & Clemens, K., 2013. From safe spaces to brave spaces. In: L.M. Landreman, ed. *The Art of Effective Facilitation: Reflections from Social Justice Educators*, 1st ed. Sterling, VA: Stylus Publishing, LLC, pp.135–50.

Arqawi, S.M., Zitawi, E.A., Rabaya, A.H., Abunasser, B.S. & Abu-Naser, S.S., 2022. Predicting university student retention using artificial intelligence. *International Journal of Advanced Computer Science and Applications*, 13(9), pp.315–24. Available at: https://doi.org/10.14569/IJACSA.2022.0130937.

Arroyo, I., Ferguson, K., Johns, J., Dragon, T., Meheranian, H., Fisher, D., Barto, A.G., Mahadevan, S. & Woolf, B.P., 2007. Repairing disengagement with non-invasive interventions. In: R. Luckin, K.R. Koedinger & J.E. Greer, eds. *Proceedings of the*

13th International Conference on Artificial Intelligence in Education (AIED 2007), Amsterdam: IOS Press. pp.195–202.

Athaluri, S., Manthena, S., Kesapragada, V.S.R.K.M., Yarlagadda, V., Dave, T. & Duddumpudi, R.T.S., 2023. Exploring the boundaries of reality: Investigating the phenomenon of artificial intelligence hallucination in scientific writing through ChatGPT references. *Cureus*, 15(4), Article e37432. Available at: https://doi.org/10.7759/cureus.37432.

Atkins, C., Girgente, G., Shirzaei, M. & Kim, J., 2024. Generative AI tools can enhance climate literacy but must be checked for biases and inaccuracies. *Communications Earth & Environment*, 5(1), Article 226. Available at: https://doi.org/10.1038/s43247-024-01392-w.

Bahroun, Z., Anane, C., Ahmed, V. & Zacca, A., 2023. Transforming education: A comprehensive review of generative artificial intelligence in educational settings through bibliometric and content analysis. *Sustainability*, 15(17), Article 12983. Available at: https://doi.org/10.3390/su151712983.

Baker, R.S., 2016. Stupid tutoring systems, intelligent humans. *International Journal of Artificial Intelligence in Education*, 26, pp.600–14. Available at: https://doi.org/10.1007/s40593-016-0105-0.

Barr, A., 2023. OpenAI just admitted it cannot identify AI-generated text. That is bad for the internet and it could be really bad for AI models. *Business Insider*, 28 July. Available at: https://www.businessinsider.com/openai-cant-identify-ai-generated-text-bad-for-internet-models-2023-7 [Accessed 17 April 2025].

Bates, A.T. & Sangra, A., 2011. *Managing technology in higher education: Strategies for transforming teaching and learning*. San Francisco, CA: Jossey-Bass.

Batra, K., Sharma, M., Batra, R., Singh, T.P. & Schvaneveldt, N., 2021. Assessing the psychological impact of COVID-19 among college students: Evidence from 15 countries. *Healthcare*, 9(2), p.222. Available at: https://doi.org/10.3390/healthcare9020222.

Bayne, S., 2015. Teacherbot: Interventions in automated teaching. *Teaching in Higher Education*, 20(4), pp.455–67. Available at: https://doi.org/10.1080/13562517.2015.1020783.

Belenguer, L., 2022. AI bias: Exploring discriminatory algorithmic decision-making models and the application of possible machine-centric solutions adapted from the pharmaceutical industry. *AI and Ethics*, 2(4), pp.771–87. Available at: https://doi.org/10.1007/s43681-022-00138-8.

Bender, E.M., 2023. Cleaning up a baby peacock sullied by a non-information spill. Available at: https://medium.com/@emilymenonbender/cleaning-up-a-baby-peacock-sullied-by-a-non-information-spill-d2e2aa642134 [Accessed 17 April 2025].

Bender, E.M., Gebru, T., McMillan-Major, A. & Shmitchell, S., 2021. On the dangers of stochastic parrots: Can language models be too big? In: *Proceedings of the 2021 ACM*

Conference on Fairness, Accountability, and Transparency, pp.610–23. Available at: https://doi.org/10.1145/3442188.3445922.

Biesta, G., 2015. What is education for? On good education, teacher judgement, and educational professionalism. *European Journal of Education*, 50(1), pp.75–87. Available at: https://doi.org/10.1111/ejed.12109.

Biggs, J., 2003. *Aligning teaching for constructing learning*. Available at: https://www.advance-he.ac.uk/knowledge-hub/aligning-teaching-constructing-learning [Accessed 17 April 2025].

Bloxham, S., den-Outer, B., Hudson, J. & Price, M., 2015. Let's stop the pretence of consistent marking: Exploring the multiple limitations of assessment criteria. *Assessment & Evaluation in Higher Education*, 41(3), pp.466–81. Available at: https://doi.org/10.1080/02602938.2015.1024607.

Bond, M., Khosravi, H., De Laat, M., Bergdahl, N., Negrea, V., Oxley, E., Pham, P., Chong, S.W. & Siemens, G., 2024. A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour. *International Journal of Educational Technology in Higher Education*, 21(1), Article 4. Available at: https://doi.org/10.1186/s41239-023-00436-z.

Bovill, C., Bulley, C.J. & Morss, K., 2011. Engaging and empowering first-year students through curriculum design: Perspectives from the literature. *Teaching in Higher Education*, 16(2), pp.197–209. Available at: https://doi.org/10.1080/13562517.2010.515024.

Bowker, L., 2021. Promoting linguistic diversity and inclusion. *The International Journal of Information, Diversity, & Inclusion*, 5(3), pp.127–51. Available at: https://doi.org/10.33137/ijidi.v5i3.36159.

Brodkin, J., 2023. Lawyers have real bad day in court after citing fake cases made up by ChatGPT. *Ars Technica*, 23 June. Available at: https://arstechnica.com/tech-policy/2023/06/lawyers-have-real-bad-day-in-court-after-citing-fake-cases-made-up-by-chatgpt/ [Accessed 17 April 2025].

Bruffee, K.A., 1999. *Collaborative learning: Higher education, interdependence, and the authority of knowledge*. Baltimore, MD: Johns Hopkins University Press.

Buçinca, Z., Malaya, M.B. & Gajos, K.Z., 2021. To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), Article 188, pp.1–21. Available at: https://doi.org/10.1145/3449287.

Campbell, P.I., 2022. 'Pray(ing) the person marking your work is not racist': Racialised inequities in HE assessment practice. *Teaching in Higher Education*, 29(5), pp.1166–80. Available at: https://doi.org/10.1080/13562517.2022.2119075.

Castaldo, J., 2023. The ABCs of AI; AI depends on a global army of poorly paid online gig workers to teach the technology what it should and should not do. *The Globe and Mail*, 16 September. Available at: https://go.gale.com/ps/i.do?id=GALE%7CA765477089&sid=googleScholar&v=2.1&it=r&linkaccess=abs&issn=03190714&p=HRC

A&sw=w&userGroupName=anon%7Eb76a97d&aty=open-web-entry [Accessed 17 April 2025].

Chan, C.K.Y., 2023. A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20, Article 38. Available at: https://doi.org/10.1186/s41239-023-00408-3.

Chatterjee, S. & Bhattacharjee, K.K., 2020. Adoption of artificial intelligence in higher education: A quantitative analysis using structural equation modelling. *Education and Information Technologies*, 25, pp.3443–63. Available at: https://doi.org/10.1007/s10639-020-10159-7.

Chatterjee, R. & Correia, A.P., 2019. Online students' attitudes toward collaborative learning and sense of community. *American Journal of Distance Education*, 34(1), pp.53–68. Available at: https://doi.org/10.1080/08923647.2020.1703479.

Chen, Y., Jensen, S., Albert, L.J., Gupta, S. & Lee, T., 2023. Artificial intelligence (AI) student assistants in the classroom: Designing chatbots to support student success. *Information Systems Frontiers*, 25(1), pp.161–82. Available at: https://doi.org/10.1007/s10796-022-10291-4.

Chew, S.L. & Cerbin, W.J., 2021. The cognitive challenges of effective teaching. *The Journal of Economic Education*, 52(1), pp.17–40. Available at: https://doi.org/10.1080/00220485.2020.1845266.

Coffield, F., Moseley, D., Hall, E. & Ecclestone, K., 2004. *Learning styles and pedagogy in post-16 learning: A systematic and critical review*. London: Learning and Skills Research Centre.

Crawford, J., Cowling, M. & Allen, K.A., 2023. Leadership is needed for ethical ChatGPT: Character, assessment, and learning using artificial intelligence (AI). *Journal of University Teaching & Learning Practice*, 20(3), Article 2. Available at: https://doi.org/10.53761/1.20.3.02.

Dale, R., 2023. Navigating the text generation revolution: Traditional data-to-text NLG companies and the rise of ChatGPT. *Natural Language Engineering*, 29(4), pp.1188–97. Available at: https://doi.org/10.1017/S1351324923000347.

D'Alfonso, S., 2020. AI in mental health. *Current Opinion in Psychology*, 36, pp.112–17. Available at: https://doi.org/10.1016/j.copsyc.2020.04.005.

Dawson, P., 2020. *Defending assessment security in a digital world: Preventing e-cheating and supporting academic integrity in higher education*. London: Routledge.

De la Harpe, B. & Radloff, A., 2000. Informed teachers and learners: The importance of assessing the characteristics needed for lifelong learning. *Studies in Continuing education*, 22(2), pp. 169–82. Available at: https://doi.org/10.1080/713695729.

Dell'Acqua, F., 2022. Falling asleep at the wheel. Available at: https://www.almendron.com/tribuna/wp-content/uploads/2023/09/falling-asleep-at-the-whee.pdf [Accessed 17 April 2025].

Dell'Acqua, F., McFowland III, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F. & Lakhani, K.R., 2023. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge

worker productivity and quality. *Harvard Business School Technology & Operations Management Unit Working Paper*, No. 24-013. *The Wharton School Research Paper*. Available at: https://ssrn.com/abstract=4573321 [Accessed 17 April 2025].

de Roock, R.S., 2024. To become an object among objects: Generative artificial 'intelligence', writing, and linguistic white supremacy. *Reading Research Quarterly*, 59(4), pp.590–608. Available at: https://doi.org/10.1002/rrq.569.

Doo, M.Y. & Bonk, C.J., 2020. The effects of self-efficacy, self-regulation and social presence on learning engagement in a large university class using flipped Learning. *Journal of Computer Assisted Learning*, 36(6), pp.997–1010. Available at: https://doi.org/10.1111/jcal.12455.

Edwards, B., 2023. Why AI writing detectors do not work: Can AI writing detectors be trusted? We dig into the theory behind them. *Ars Technica*, 14 July. Available at: https://arstechnica.com/information-technology/2023/07/why-ai-detectors-think-the-us-constitution-was-written-by-ai/ [Accessed 17 April 2025].

European Union, 2024. Regulation (EU) 2024/1689 of the European Parliament and of the council of 13 June 2024 laying down harmonised rules on artificial intelligence. Available at: http://data.europa.eu/eli/reg/2024/1689/oj [Accessed 17 April 2025].

Farrelly, T. & Baker, N., 2023. Generative artificial intelligence: Implications and considerations for higher education practice. *Education Sciences*, 13(11), Article 1109. Available at: https://doi.org/10.3390/educsci13111109.

Ferrara, E., 2023. Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies. *Sci*, 6(1), Article 3. Available at: https://doi.org/10.3390/sci6010003.

Ferrer, X., van Nuenen, T., Such, J.M., Coté, M. & Criado, N., 2021. Bias and discrimination in AI: A cross-disciplinary perspective. *IEEE Technology and Society Magazine*, 40(2), pp.72–80. Available at: https://doi.org/10.1109/MTS.2021.3056293.

Firat, M., 2023. What ChatGPT means for universities: Perceptions of scholars and students. *Journal of Applied Learning and Teaching*, 6(1), pp.57–63. Available at: https://doi.org/10.37074/jalt.2023.6.1.22.

Fitzpatrick, K.K., Darcy, A. & Vierhile, M., 2017. Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. *JMIR Mental Health*, 4(2), Article e19. Available at: https://doi.org/10.2196/mental.7785.

Fleckenstein, J., Meyer, J., Jansen, T., Keller, S.D., Köller, O. & Möller, J., 2024. Do teachers spot AI? Evaluating the detectability of AI-generated texts among student essays. *Computers and Education: Artificial Intelligence*, 6, Article 100209. Available at: https://doi.org/10.1016/j.caeai.2024.100209.

Floridi, L., 2023. AI as agency without intelligence: On ChatGPT, large language models, and other generative models. *Philosophy & Technology*, 36, Article 15. Available at: https://doi.org/10.1007/s13347-023-00621-y.

Forsyth, R., 2022. *Confident assessment in higher education*. London: Sage.

Forsyth, R., Hamshire, C., Fontaine-Rainen, D. & Soldaat, L., 2022. Shape-shifting and pushing against the odds: Staff perceptions of the experiences of first-generation students in South Africa and the UK. *The Australian Educational Researcher*, 49(2), pp.307–21. Available at: https://doi.org/10.1007/s13384-021-00438-8.

Fredricks, J.A., Reschly, A.L. & Christenson, S.L. eds., 2019. *Handbook of student engagement interventions: Working with disengaged students*. Cambridge, MA: Academic Press.

Fullan, M. & Langworthy, M. (2014). *A rich seam: How new pedagogies find deep learning*. New Zealand: OER4Teachers OER Collection. Available at: https://staging.oer4pacific.org/id/eprint/5/.

Garcia, D., 2023. Influence of Facebook algorithms on political polarization tested. *Nature*, 620, pp.39–41. Available at: https://doi.org/10.1038/d41586-023-02325-x.

Gasevic, D., Tsai, Y.S., Dawson, S. & Pardo, A., 2019. How do we start? An approach to learning analytics adoption in higher education. *The International Journal of Information and Learning Technology*, 36(4), pp.342–53. Available at: https://doi.org/10.1108/IJILT-02-2019-0024.

Goetze, T.S., 2024. AI art is theft: Labour, extraction, and exploitation: Or, on the dangers of stochastic Pollocks. In: *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, 3–6 June, Rio de Janeiro, Brazil. New York, NY: ACM, pp.186–96. Available at: https://doi.org/10.1145/3630106.3658898.

Graesser, A.C., Hu, X., Rus, V. & Cai, Z., 2020. Conversation-based learning and assessment environments. In: M.D. Shermis & J. Burstein (eds.) *Handbook of Automated Scoring*. 1st ed. Boca Raton, FL: Chapman and Hall/CRC, pp.20. Available at: https://doi.org/10.1201/9781351264808.

Graesser, A.C., Wiemer-Hastings, P., Wiemer-Hastings, K., Harter, D., Tutoring Research Group, T.R.G. & Person, N., 2000. Using latent semantic analysis to evaluate the contributions of students in AutoTutor. *Interactive Learning Environments*, 8(2), pp.129–47. Available at: DOI: 10.1076/1049-4820(200008)8:2;1-B;FT129.

Graham, K., 2018. TechMatters: Doing digital curation with Wakelet. *LOEX Quarterly*, 45(2), pp.4–6.

Gravett, K. & Winstone, N.E., 2022. Making connections: Authenticity and alienation within students' relationships in higher education. *Higher Education Research & Development*, 41(2), pp.360–74. Available at: https://doi.org/10.1080/07294360.2020.1842335.

Grosvenor, B., 2013. Bad mouthing Lely's hands. *Art History News*. Available at: https://www.arthistorynews.com/articles/2045_Bad_mouthing_Lelys_hands [Accessed 17 April 2025].

Guleria, P. & Sood, M., 2023. Explainable AI and machine learning: Performance evaluation and explainability of classifiers on educational data mining inspired career counseling. *Education and Information Technologies*, 28(1), pp.1081–116. Available at: https://doi.org/10.1007/s10639-022-11221-2.

Gunawardena, A., 2017. Brief survey of analytics in K12 and higher education. *International Journal on Innovations in Online Education*, 1(1). Available at: https://doi.org/10.1615/IntJInnovOnlineEdu.v1.i1.80.

Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S. & Yang, G.Z., 2019. XAI—Explainable artificial intelligence. *Science Robotics*, 4(37), Article eaay7120. Available at: https://doi.org/10.1126/scirobotics.aay7120.

Guraya, S.Y., 2023. The bandwagon of artificial intelligence and its untapped areas in medical education and scholarly research. *Advances in Biomedical and Health Sciences*, 2(3), pp.103–4. Available at: https://doi.org/10.4103/abhs.abhs_42_23.

Hall, G.C.N., Berkman, E.T., Zane, N.W., Leong, F.T., Hwang, W.C., Nezu, A.M., Nezu, C.M., Hong, J.J., Chu, J.P. & Huang, E.R., 2021. Reducing mental health disparities by increasing the personal relevance of interventions. *American Psychologist*, 76(1), p.91. Available at: https://psycnet.apa.org/doi/10.1037/amp0000616.

Hanley, T., 2021. Researching online counselling and psychotherapy: The past, the present and the future. *Counselling and Psychotherapy Research*, 21(3), pp.493–7. Available at: https://doi.org/10.1002/capr.12385.

Henderson, M., Selwyn, N. & Aston, R., 2015. What works and why? Student perceptions of 'useful' digital technology in university teaching and learning. *Studies in Higher Education*, 42(8), pp.1567–79. Available at: https://doi.org/10.1080/03075079.2015.1007946.

Hinds, J., Williams, E.J. & Joinson, A.N., 2020. It would not happen to me: Privacy concerns and perspectives following the Cambridge Analytica scandal. *International Journal of Human-Computer Studies*, 143, Article 102498. Available at: https://doi.org/10.1016/j.ijhcs.2020.102498.

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S.B., Santos, O.C., Rodrigo, M.T., Cukurova, M., Bittencourt, I.I. & Koedinger, K.R., 2022. Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32, pp.504–26. Available at: https://doi.org/10.1007/s40593-021-00239-1.

Holmes, W. & Tuomi, I., 2022. State of the art and practice in AI in education. *European Journal of Education*, 57(4), pp.542–70. Available at: https://doi.org/10.1111/ejed.12533.

Hosseini, M., Gao, P. & Vivas-Valencia, C., 2025. A social-environmental impact perspective of generative artificial intelligence. *Environmental Science and Ecotechnology*, 23, Article 100520. Available at: https://doi.org/10.1016/j.ese.2024.100520.

Hsieh, J.P.A., Rai, A. & Keil, M., 2008. Understanding digital inequality: Comparing continued use behavioral models of the socio-economically advantaged and disadvantaged. *MIS Quarterly*, 32(1), pp.97–126. Available at: https://doi.org/10.2307/25148830.

Hsu, T., 2023. What can you do when A.I. lies about you? *The New York Times*, 3 August. Available at: https://www.nytimes.com/2023/08/03/business/media/ai-defamation-lies-accuracy.html?smid=tw-share&s=03 [Accessed 17 April 2025].

Huggins-Manley, A.C., Booth, B.M. & D'Mello, S.K., 2022. Toward argument-based fairness with an application to AI-enhanced educational assessments. *Journal of*

Educational Measurement, 59(3), pp.362–88. Available at: https://doi.org/10.1111/jedm.12334.

Hussain, M., Zhu, W., Zhang, W. & Abidi, S.M.R., 2018. Student engagement predictions in an e-learning system and their impact on student course assessment scores. *Computational Intelligence and Neuroscience*. Available at: https://doi.org/10.1155/2018/6347186.

Ihantola, P., Vihavainen, A., Ahadi, A., Butler, M., Börstler, J., Edwards, S.H., Isohanni, E., Korhonen, A., Petersen, A., Rivers, K. & Rubio, M.Á., 2015. Educational data mining and learning analytics in programming: Literature review and case studies. *Proceedings of the 2015 ITiCSE on Working Group Reports*, pp.41–63. Available at: https://doi.org/10.1145/2858796.2858798.

Ilić, M.P., Păun, D., Popović Šević, N., Hadžić, A. & Jianu, A., 2021. Needs and performance analysis for changes in higher education and implementation of artificial intelligence, machine learning, and extended reality. *Education Sciences*, 11(10), p.568. Available at: https://doi.org/10.3390/educsci11100568.

Illingworth, S. & Radhakrishnan, M.L., 2023. 'I am here because I wanted to shine': How poetry can be used to better understand undergraduate students' first-year chemistry or related course experiences. *Chemistry Education Research and Practice*. Available at: https://doi.org/10.1039/D2RP00276K.

Jandrić, P. & Ford, D.R., 2022. Postdigital ecopedagogies: Genealogies, contradictions, and possible futures. *Postdigit Sci Educ*, 4, pp.692–710. Available at: https://doi.org/10.1007/s42438-020-00207-3.

Jokhan, A., Sharma, B. & Singh, S., 2019. Early warning system as a predictor for student performance in higher education blended courses. *Studies in Higher Education*, 44(11), pp.1900–11. Available at: https://doi.org/10.1080/03075079.2018.1466872.

Kaplan-Rakowski, R., Grotewold, K., Hartwick, P. & Papin, K., 2023. Generative AI and teachers' perspectives on its implementation in education. *Journal of Interactive Learning Research*, 34(2), pp.313–38. Waynesville, NC: Association for the Advancement of Computing in Education (AACE). Available at: https://www.learntechlib.org/primary/p/222363/ [Accessed 17 April 2025].

Karpouzis, K., Pantazatos, D., Taouki, J. & Meli, K., 2024. Tailoring education with GenAI: A new horizon in lesson planning. *arXiv preprint*, arXiv:2403.12071. Available at: https://doi.org/10.48550/arXiv.2403.12071.

Kharbat, F.F., Alshawabkeh, A. & Woolsey, M.L., 2021. Identifying gaps in using artificial intelligence to support students with intellectual disabilities from education and health perspectives. *Aslib Journal of Information Management*, 73(1), pp.101–28. Available at: https://doi.org/10.1108/AJIM-02-2020-0054.

Kietzmann, J., Lee, L.W., McCarthy, I.P. & Kietzmann, T.C., 2020. Deepfakes: Trick or treat? *Business Horizons*, 63(2), pp.135–46. Available at: https://doi.org/10.1016/j.bushor.2019.11.006.

Kim, J., Merrill, K., Xu, K. & Sellnow, D.D., 2020. My teacher is a machine: Understanding students' perceptions of AI teaching assistants in online education.

International Journal of Human–Computer Interaction, 36(20), pp.1902–11. Available at: https://doi.org/10.1080/10447318.2020.1801227.

Koehler, M. & Mishra, P., 2009. What is technological pedagogical content knowledge (TPACK)? *Contemporary Issues in Technology and Teacher Education*, 9(1), pp.60–70. Available at: https://citejournal.org/volume-9/issue-1-09/general/what-is-technological-pedagogicalcontent-knowledge [Accessed 17 April 2025].

Koutsouleris, Nikolaos, Hauser, Tobias U., Skvortsova, Vasilisa & De Choudhury, Munmun, 2022. From promise to practice: Towards the realisation of AI-informed mental health care. *The Lancet Digital Health,* 4(11), pp.e829–40. Available at: DOI: 10.1016/S2589-7500(22)00153-4.

Kovanović, V., Joksimović, S., Poquet, O., Hennis, T., Čukić, I., De Vries, P., Hatala, M., Dawson, S., Siemens, G. & Gašević, D., 2018. Exploring communities of inquiry in massive open online courses. *Computers & Education*, 119, pp.44–58. Available at: https://doi.org/10.1016/j.compedu.2017.11.010.

Kretzschmar, K., Tyroll, H., Pavarini, G., Manzini, A., Singh, I. & NeurOx Young People's Advisory Group, 2019. Can your phone be your therapist? Young people's ethical perspectives on the use of fully automated conversational agents (chatbots) in mental health support. *Biomedical Informatics Insights*, 11. Available at: https://doi.org/10.1177/1178222619829083.

Kulik, J.A. & Fletcher, J.D., 2016. Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), pp.42–78. Available at: https://doi.org/10.3102/0034654315581420.

Kumar, D., Kelley, P.G., Consolvo, S., Mason, J., Bursztein, E., Durumeric, Z., Thomas, K. & Bailey, M., 2021. Designing toxic content classification for a diversity of perspectives. *arXiv preprint*, arXiv:2106.04511. Available at: https://doi.org/10.48550/arXiv.2106.04511.

Lee, K.F. & Qiufan, C., 2021. *AI 2041: Ten visions for our future*. New York, NY: Crown Publishing Group.

Lee, T., Jagannath, K., Aggarwal, N., Sridar, R., Wilde, S., Hill, T. & Chen, Y., 2019. Intelligent career advisers in your pocket? A need assessment study of chatbots for student career advising. *AMCIS 2019 Proceedings*. Available at: https://aisel.aisnet.org/amcis2019/human_computer_interact/human_computer_interact/14.

Lee, T. B. & Trott, S., 2023. Large language models, explained with a minimum of math and jargon. *Understanding AI*. Available at: https://www.understandingai.org/p/large-language-models-explained-with?s=03.

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E. & Zou, J., 2023. GPT detectors are biased against non-native English writers. *Patterns*, 4(7), Article 100779. Available at: https://doi.org/10.1016/j.patter.2023.100779.

Lipson, S.K. & Sonneville, K.R., 2020. Understanding suicide risk and eating disorders in college student populations: Results from a National Study. *International Journal of Eating Disorders*, 53(2), pp.229–38. Available at: https://doi.org/10.1002/eat.23188.

Liu, L., 2024. Stanford misinformation expert accused of using AI to fabricate court statement. *The Stanford Daily*. Available at: https://stanforddaily.com/2024/12/02/jeff-hancock-court-declaration/.

Lloyd, A. & Forsyth, R., 2024. The role of internal and external moderation for assuring academic standards. In: N. Reimann, I. Sadler & J. Hill, eds. *Academic Standards in Higher Education: Critical Perspectives and Practical Strategies*. London: Routledge, pp.57–70.

Loebner, H., 2009. How to hold a Turing test contest. In: R. Epstein, G. Roberts & G. Beber, eds. *Parsing the Turing Test*. Dordrecht: Springer, pp.173–9.

Luckin, R., Holmes, W., Griffiths, M. & Forcier, L.B., 2016. Intelligence unleashed: An Argument for AI in Education. *Pearson*. Available at: https://www.pearson.com/content/dam/corporate/global/pearson-dot-com/files/innovation/Intelligence-Unleashed-Publication.pdf [Accessed 17 April 2025].

Luo, J., 2024. A critical review of GenAI policies in higher education assessment: A call to reconsider the 'originality' of students' work. *Assessment & Evaluation in Higher Education*, 49(5), pp.651–64. Available at: https://doi.org/10.1080/02602938.2024.2309963.

Ma, Z., Mei, Y. & Su, Z., 2023. Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support. AMIA Annu Symp Proc. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC10785945/.

Malmström, H., Stöhr, C. & Ou, A.W., 2023. Chatbots and other AI for learning: A survey of use and views among university students in Sweden. *Chalmers Studies in Communication and Learning in Higher Education*, 2023, p.1. Available at: https://doi.org/10.17196/cls.csclhe/2023/0 [Accessed 17 April 2025].

McArthur, J., 2023. Rethinking authentic assessment: Work, well-being, and society. *Higher Education*, 85, pp.85–101. Available at: https://doi.org/10.1007/s10734-022-00822-y.

McCormack, M., 2023. EDUCAUSE QuickPoll results: Adopting and adapting to generative AI in higher ed tech. *EDUCAUSE*. Available at: https://er.educause.edu/articles/2023/4/educause-quickpoll-results-adopting-and-adapting-to-generative-ai-in-higher-ed-tech [Accessed 17 April 2025].

Mhlanga, D., 2023. Open AI in education, the responsible and ethical use of ChatGPT towards lifelong learning. In: H.S. Sætra, ed. *FinTech and Artificial Intelligence for Sustainable Development: The Role of Smart Technologies in Achieving Development Goals*. Cham: Palgrave Macmillan, pp.387–409. Available at: https://doi.org/10.1007/978-3-031-37776-1_17.

Mitrovic, A., 2003. An intelligent SQL tutor on the web. *International Journal of Artificial Intelligence in Education*, 13(2–4), pp.173–97.

Morgenstern, J.D., Rosella, L.C., Daley, M.J., Goel, V., Schünemann, H.J. & Piggott, T., 2021. 'AI's gonna have an impact on everything in society, so it has to have an impact on public health': A fundamental qualitative descriptive study of the implications of

artificial intelligence for public health. *BMC Public Health*, 21, Article 40. Available at: https://doi.org/10.1186/s12889-020-10030-x.

Morreale, F., Bahmanteymouri, E., Burmester, B., Chen, A., & Thorp, M. (2023). The unwitting labourer: extracting humanness in AI training. *AI & SOCIETY*. doi:10.1007/s00146-023-01692-3

Morreale, F., Bahmanteymouri, E., Burmester, B., Chen, A. & Thorp, M., 2024. The unwitting labourer: Extracting humanness in AI training. *AI & Society*, 39, pp.2389–99. Available at: https://doi.org/10.1007/s00146-023-01692-3.

Morris, M.R., Fourney, A., Ali, A. & Vonessen, L., 2018. Understanding the needs of searchers with dyslexia. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*, Montreal, QC, Canada. New York, NY: Association for Computing Machinery, pp.1–12. Available at: https://doi.org/10.1145/3173574.3173609.

Murugesan, S. & Cherukuri, A.K., 2023. The rise of generative artificial intelligence and its impact on education: The promises and perils. *Computer*, 56(5), pp.116–21. Available at: https://doi.org/10.1109/MC.2023.3253292.

Narayanan, A. & Kapoor, S., 2024. *AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference*. Princeton, NJ: Princeton University Press.

Nature, 2023. Tools such as ChatGPT threaten transparent science; here are our ground rules for their use. *Nature*, 613, p.612. Available at: https://doi.org/10.1038/d41586-023-00191-1.

Ng, D.T.K., Lee, M., Tan, R.J.Y., Hu, X., Downie, J.S. & Chu, S.K.W., 2023. A review of AI teaching and learning from 2000 to 2020. *Education and Information Technologies*, 28, pp.8445–501. Available at: https://doi.org/10.1007/s10639-022-11491-w.

Nguyen, A., Ngo, H.N., Hong, Y., Dang, B. & Nguyen, B.P.T., 2023. Ethical principles for artificial intelligence in education. *Education and Information Technologies*, 28, pp.4221–41. Available at: https://doi.org/10.1007/s10639-022-11316-w.

Nieminen, J.H., 2022. Assessment for inclusion: Rethinking inclusive assessment in higher education. *Teaching in Higher Education*, 29(4), pp.841–59. Available at: https://doi.org/10.1080/13562517.2021.2021395.

Nieminen, J.H. & Lahdenperä, J., 2021. Assessment and epistemic (in)justice: How assessment produces knowledge and knowers. *Teaching in Higher Education*, 29(1), pp.300–17. Available at: https://doi.org/10.1080/13562517.2021.1973413.

Noy, S. & Zhang, W., 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), pp.187–92. Available at: https://doi.org/10.1126/science.adh2586.

Oprean, D. & Balakrishnan, B., 2020. From engagement to user experience: A theoretical perspective towards immersive learning. In: M. Schmidt, A.A. Tawfik, I. Jahnke & Y. Earnshaw, eds. *Learner and User Experience Research: An Introduction for the Field of Learning Design & Technology*. EdTech Books, pp.199–216. Available

at: https://edtechbooks.org/ux/10_from_engagement_to_user_experience [Accessed 17 April 2025].

Ouyang, F., Zheng, L. & Jiao, P., 2022. Artificial intelligence in online higher education: A systematic review of empirical research from 2011 to 2020. *Education and Information Technologies*, 27, pp.7893–925. Available at: https://doi.org/10.1007/s10639-022-10925-9.

Owusu-Ansah, F.E., Addae, A.A., Peasah, B.O., Oppong Asante, K. & Osafo, J., 2020. Suicide among university students: Prevalence, risks and protective factors. *Health psychology and behavioral medicine*, 8(1), pp.220–33. Available at: https://doi.org/10.1080/21642850.2020.1766978.

Panadero, E., Jonsson, A. & Botella, J., 2017. Effects of self-assessment on self-regulated learning and self-efficacy: Four meta-analyses. *Educational Research Review*, 22, pp.74–98. Available at: https://doi.org/10.1016/j.edurev.2017.08.004.

Pedler, M.L., Willis, R. & Nieuwoudt, J.E., 2022. A sense of belonging at university: Student retention, motivation and enjoyment. *Journal of Further and Higher Education*, 46(3), pp.397–408. Available at: https://doi.org/10.1080/030987 7X.2021.1955844.

Perkins, M., Furze, L., Roe, J. & MacVaugh, J., 2024. The artificial intelligence assessment scale (AIAS): A framework for ethical integration of generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, 21(6): 49–66.

Plata, S., De Guzman, M.A. & Quesada, A., 2023. Emerging research and policy themes on academic integrity in the age of chat GPT and generative AI. *Asian Journal of University Education*, 19(4), pp.743–58. Available at: https://doi.org/10.24191/ajue.v19i4.24697.

Poort, I., Jansen, E. & Hofman, A., 2022. Does the group matter? Effects of trust, cultural diversity, and group formation on engagement in group work in higher education. *Higher Education Research & Development*, 41(2), pp.511–26. Available at: https://doi.org/10.1080/07294360.2020.1839024.

Poritz, I., 2024. OpenAI fails to escape first defamation suit from radio host. *Bloomberg Law*, 16 January. Available at: https://news.bloomberglaw.com/ip-law/openai-fails-to-escape-first-defamation-suit-from-radio-host [Accessed 17 April 2025].

Reich, J. & Ruipérez-Valiente, J.A., 2019. The MOOC pivot. *Science*, 363(6423), pp.130–1. Available at: https://doi.org/10.1126/science.aav7958.

Rillig, M.C., Ågerstrand, M., Bi, M., Gould, K.A. & Sauerland, U., 2023. Risks and benefits of large language models for the environment. *Environmental Science & Technology*, 57(9), pp.3464–6. Available at: https://doi.org/10.1021/acs.est.3c01106.

Roe, J. & Perkins, M., 2024. Deepfakes and higher education: A research agenda and scoping review of synthetic media. *Journal of University Teaching and Learning Practice*, 21(10). Available at: https://doi.org/10.53761/2y2np178.

Romero, C. & Ventura, S., 2020. Educational data mining and learning analytics: An updated survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(3), Article e1355. Available at: https://doi.org/10.1002/widm.1355.

Rowe, N., 2023. The human toll of training AI: Content moderators in Kenya faced trauma for $2 an hour. *The Guardian*, 2 August. Available at: https://www.theguardian.com/technology/2023/aug/02/ai-chatbot-training-human-toll-content-moderator-meta-openai [Accessed 17 April 2025].

Rozado, D., 2023. The political biases of ChatGPT. *Social Sciences*, 12(3), Article 148. Available at: https://doi.org/10.3390/socsci12030148.

Rumsfeld, D., 2002. Press Briefing. *US Department of Defense*, 12 February. Available at: https://usinfo.org/wf-archive/2002/020212/epf202.htm [Accessed 17 April 2025].

Samuelson, P., 2023. Generative AI meets copyright: Ongoing lawsuits could affect everyone who uses generative AI. *Science*, 381(6654), pp.158–61. Available at: https://doi.org/10.1126/science.adi0656.

Sarı, T., Nayir, F. & Bozkurt, A., 2024. Reimagining education: Bridging artificial intelligence, transhumanism, and critical pedagogy. *Journal of Educational Technology and Online Learning*, 7(1), pp.102–15. http://doi.org/10.31681/jetol.1308022.

Selwyn, N., 2019a. *Should robots replace teachers?: AI and the future of education*. Cambridge: Polity Press.

Selwyn, N., 2019b. What is the problem with learning analytics? *Journal of Learning Analytics*, 6(3), pp.11–19. Available at: https://doi.org/10.18608/jla.2019.63.3.

Selwyn, N., 2023. Lessons to be learnt? Education, techno-solutionism, and sustainable development. In: H.S. Sætra, ed. *Technology and Sustainable Development: The Promise and Pitfalls of Techno-Solutionism*. New York, NY: Routledge, pp.71–83.

Sghir, N., Adadi, A. & Lahmer, M., 2023. Recent advances in predictive learning analytics: A decade systematic review (2012–2022). *Education and Information Technologies*, 28, pp.8299–333. Available at: https://doi.org/10.1007/s10639-022-11536-0.

Sharma, V., Goyal, M. & Malik, D., 2017. An intelligent behaviour shown by chatbot system. *International Journal of New Technology and Research*, 3(4), pp.52–4.

Siemens, G., 2013. Learning analytics: The emergence of a discipline. *American Behavioral Scientist*, 57(10), pp.1380–400. Available at: https://doi.org/10.1177/0002764213498851.

Silva, H.E.C. da, Santos, G.N.M., Leite, A.F., Mesquita, C.R.M., Figueiredo, P.T. da S., Stefani, C.M. & de Melo, N.S., 2023. The use of artificial intelligence tools in cancer detection compared to the traditional diagnostic imaging methods: An overview of the systematic reviews. *PLOS ONE*, 18(10), Article e0292063. Available at: https://doi.org/10.1371/journal.pone.0292063.

Southworth, J., Migliaccio, K., Glover, J., Reed, D., McCarty, C., Brendemuhl, J. & Thomas, A., 2023. Developing a model for AI across the curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*, 4, Article 100127. Available at: https://doi.org/10.1016/j.caeai.2023.100127.

Sternberg, R.J., 2024. Human intelligence. *Encyclopedia Britannica*, 5 April. Available at: https://www.britannica.com/science/human-intelligence-psychology [Accessed 17 April 2025].

Sterne, J. & Sawhney, M., 2022. The acousmatic question and the will to datafy: Otter.ai, low-resource languages, and the politics of machine listening. *Kalfou*, 9(2), pp.288–306. Available at: https://doi.org/10.15367/kf.v9i2.617.

Strayhorn, T.L., 2018. *College students' sense of belonging: A key to educational success for all students*. New York: Routledge.

Struckman, I. & Kupiec, S., 2023. Why they are worried: Examining experts' motivations for signing the 'Pause Letter'. *arXiv preprint*, arXiv:2306.00891. Available at: https://doi.org/10.48550/arXiv.2306.00891.

Tai, J., Ajjawi, R. & Umarova, A., 2021. How do students experience inclusive assessment? A critical review of contemporary literature. *International Journal of Inclusive Education*, 28(9), pp.1936–53. Available at: https://doi.org/10.1080/13603116.2021.2011441.

Thomas, L. (2012). *Building student engagement and belonging in Higher Education at a time of change: Final report from the What Works? Student retention & success programme*. York: Paul Hamlyn Foundation.

Tice, D., Baumeister, R., Crawford, J., Allen, K.A. & Percy, A., 2021. Student belongingness in higher education: Lessons for Professors from the COVID-19 pandemic. *Journal of University Teaching & Learning Practice*, 18(4), Article 2. Available at: https://doi.org/10.53761/1.18.4.2.

Timmons, A.C., Duong, J.B., Simo Fiallo, N., Lee, T., Vo, H.P.Q., Ahle, M.W., Comer, J.S., Brewer, L.C., Frazier, S.L. & Chaspari, T., 2023. A call to action on assessing and mitigating bias in artificial intelligence applications for mental health. *Perspectives on Psychological Science*, 18(5), pp.1062–96. Available at: https://doi.org/10.1177/17456916221134490.

Vaughn, M., 2020. What is student agency and why is it needed now more than ever? *Theory into Practice*, 59(2), pp.109–18. Available at: https://doi.org/10.1080/00405841.2019.1702393.

Vignesh, S., Priyanka, C.S., Manju, H.S. & Mythili, K., 2021, March. An intelligent career guidance system using machine learning. In: *2021 7th International Conference on Advanced Computing and Communication Systems (ICACCS)* (Vol. 1, pp.987–90). IEEE. Available at: https://doi.org/10.1109/ICACCS51430.2021.9441978.

Vodenko, K.V. & Lyausheva, S.A., 2020. Science and education in the form 4.0: Public policy and organization based on human and artificial intellectual capital. *Journal of Intellectual Capital*, 21(4), pp.549–64. Available at: https://doi.org/10.1108/JIC-11-2019-0274.

Waghid, F., 2019. Towards decolonisation within university education: On the innovative application of educational technology. In: C.H. Manthalu & Y. Waghid, eds. *Education for Decoloniality and Decolonisation in Africa*. Cham: Palgrave Macmillan, pp.139–53. Available at: https://doi.org/10.1007/978-3-030-15689-3_8.

Wang, Y., 2023. Artificial intelligence technologies in college English translation teaching. *Journal of Psycholinguistic Research*, 52, pp.1525–44. Available at: https://doi.org/10.1007/s10936-023-09960-5.

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P. & Waddington, L., 2023. Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19(1), Article 26. Available at: https://doi.org/10.1007/s40979-023-00146-z.

Williams, A., Miceli, M. & Gebru, T., 2022. The exploited labor behind artificial intelligence: Supporting transnational worker organizing should be at the center of the fight for 'ethical AI'. *Noema Magazine*, 13 October. Available at: https://www.noemamag.com/the-exploited-labor-behind-artificial-intelligence/ [Accessed 17 April 2025].

Williamson, B., Bayne, S. & Shay, S., 2020. The datafication of teaching in higher education: Critical issues and perspectives. *Teaching in Higher Education*, 25(4), pp.351–65. Available at: https://doi.org/10.1080/13562517.2020.1748811.

Winans, M.D., 2021. Grammarly's tone detector: Helping students write pragmatically appropriate texts. *RELC Journal*, 52(2), pp.348–52. Available at: https://doi.org/10.1177/00336882211010506.

Yamashita, I., et al., 2021. Measuring the AI content of government-funded R&D projects: A proof of concept for the OECD Fundstat initiative. *OECD Science, Technology and Industry Working Papers*, No. 2021/09. Paris: OECD Publishing. Available at: https://doi.org/10.1787/7b43b038-en.

Yang, J., 2022. Perceptions of preservice teachers on AI chatbots in English education. *International Journal of Internet, Broadcasting and Communication*, 14(1), pp.44–52. Available at: https://doi.org/10.7236/IJIBC.2022.14.1.44.

Zawacki-Richter, O., Marín, V.I., Bond, M. & Gouverneur, F., 2019. Systematic review of research on artificial intelligence applications in higher education – where are the educators? *International Journal of Educational Technology in Higher Education*, 16, Article 39. Available at: https://doi.org/10.1186/s41239-019-0171-0.

Zhai, X., 2022. ChatGPT user experience: Implications for education. *SSRN 4312418*. Available at: https://dx.doi.org/10.2139/ssrn.4312418.

Index