

Contents

Background 2

Part I: The challenges of generative AI in (BETA) education 3

- 1. *What is generative AI?* 3**
- 2. *Challenges of GenAI* 4**
- 3. *Challenges of GenAI in higher education* 5**
- 4. *What do students need to learn about GenAI?* 7**

Background

The faculty board of the VU Faculty of Science (hereinafter referred to as BETA) is aware that there is a need for clear guidelines on the use of generative AI (hereinafter referred to as GenAI) in education. VU already has several guidelines in place^{1 2}. However, BETA teachers, curriculum designers and examination committees asked for additional and/or more specific faculty guidelines, that they can fall back on. The following document has been compiled using a number of available documents and guidelines as inspiration. Sources include (non-exhaustive): VU guidelines (both VU-wide and other faculties), guidelines used by other universities and procedures used by various programmes within BETA. This document will be updated as new VU guidelines are published and as the state-of-the-art progresses. Parts of these guidelines will be incorporated within the BETA assessment policy where appropriate, as the two fields occasionally overlap.

The BETA guideline on use of GenAI consists of 3 parts:

- **Part I:** An exploration of the challenges of the use of GenAI in education and a justification of the choices BETA makes in the use of generative AI. This document serves as a context for the BETA policy and is available for all interested parties.
This document contains part I. The other parts can be accessed via <weblink>.
- **Part II:** BETA policy on the use of GenAI. This document states what is and is not allowed concerning the use of generative AI in BETA education (necessarily non-exhaustive).
- **Part III:** Guidelines on BETA curriculum design regarding learning paths for AI Literacy. In this part we give some background for programme management and teachers on how to design a curriculum aligned with the BETA policy for use of GenAI, on curriculum-, course- and assessment level.
- **Part IV:** Tips for teachers and examiners on the use of GenAI in education. This part includes some tips and best practices which have previously been shared by other sources, for example the Center for Teaching and Learning (VU CTL).

This document is a living document, since GenAI is continuously evolving, it would be presumptuous to define a definitive guideline. Both derivative versions (e.g. web-based) and this document will be updated to reflect new developments, with major changes communicated through newsletters and stakeholder emails.

Any input regarding the guidelines can be shared via <weblink>, so they can be integrated in a new iteration of the guidelines.

¹ <https://vu.nl/en/student/examinations/generative-ai-your-use-our-expectations>

² <https://vu.nl/en/student/examinations/generative-ai-your-use-our-expectations>

Part I: The challenges of generative AI in (BETA) education

Following the need for clear guidelines on the use of GenAI for BETA VU, the BETA Faculty Board defines a set of three basic principles that apply to all programmes within BETA VU:

1. Shift of focus from outcome (grade) to the process of learning.
2. GenAI literacy as an integral part of each study programme, covering both the use of GenAI and ethical aspects (both moral and environmental).
3. Clear guidelines as to what constitutes fraud and what does not.

These principles are more clearly defined in part II: BETA policy on the use of GenAI. The justification of how these principles came to be by outlining the challenges GenAI poses as a whole and to higher education in particular, is contained in this part I. This document therefore serves as a context for the BETA policy and is available for all interested parties.

1. What is generative AI?

Generative AI refers to advanced machine learning models that can create new content, such as text, images or code, by learning from existing patterns in large amounts of data, typically in the form of 'machine learning'. Despite the sudden hype in recent years, GenAI has been around since the mid-2010s, with significant advancements over the years. The release of OpenAI's ChatGPT, based on the GPT-3.5 Large Language Model in November 2022, surprised many across society by combining advanced text generation with an intuitive chat-based interface, making sophisticated AI tools broadly accessible. See, for example, the following example interaction with OpenAI's ChatGPT, based on a prompt asking information about GenAI in higher education:

What is Generative AI?

Generative AI is a type of artificial intelligence that creates new content, like text, images, music, or videos. Instead of just analyzing existing data, it learns from large amounts of information and then produces original content. For example, it can write essays, generate artwork, or compose music based on patterns it has learned.

Impact on Higher Education

Generative AI is transforming higher education in several ways:

1. *Personalized Learning: Generative AI acts as a personal tutor, offering customized explanations, practice problems, and feedback. This helps students learn at their own pace and according to their individual needs.*
2. *Educator Support: AI assists educators by creating course materials such as quizzes and assignments. It can also help with grading, allowing teachers more time to focus on interacting with students and improving their teaching methods.*
3. *Creativity Boost: In creative subjects like writing, art, or music, generative AI provides new tools for brainstorming and exploring different styles, expanding opportunities for creative learning.*
4. *Research Assistance: AI helps researchers by generating ideas, analyzing data, and summarizing large amounts of information, making research faster and more efficient.*
5. *Ethical Challenges: The use of AI in assignments raises concerns about plagiarism and academic honesty, pushing universities to rethink how they assess student work.*

In summary, generative AI is reshaping higher education by enhancing learning, supporting educators, fostering creativity, advancing research, and prompting important discussions about ethics and academic integrity.

The above text was generated using ChatGPT via <https://openai.com/chatgpt/>, using the prompt: "Explain in simple language what generative AI is, and how generative AI impacts higher education. Use no more than 250 words." (ChatGPT-4 (August 2024 version))

2. Challenges of GenAI

Please note that the above example shows that ChatGPT generated overwhelmingly positive output to a neutral prompt. This imbalance demonstrates one of the risks present in the use of such GenAI tools: output can be biased. The following list contains some of the more problematic elements of GenAI tools that are occasionally glossed over:

- **Bias**
 - **Bias in quality:** GenAI tools are based on large volumes of input data, which are compiled from a variety of sources, making it difficult to establish whether output is trustworthy. GenAI tools will present grammatical phrases that come across as competent writing, but the contents can be derived (partly) from unverified sources, marketing material or even misinformation. As such tools will typically not offer references to their sources, users have no way of verifying the information a GenAI tool offers (other than manually checking using other methods).
 - **Model bias:** Biases can occur accidentally or be deliberate. Accidental bias may occur when there is an unintentional bias in the data used to train the GenAI tool. GenAI tools can also have a deliberate bias, in which the GenAI tool is specifically designed or willfully trained to reflect that bias. Both are hard to recognize, especially because many tools do not share their source code and, if they do, most users are not trained to read the code.
 - **Hallucinations:** GenAI tools ‘hallucinate’ when they produce outputs that sound plausible but are factually incorrect or fabricated. GenAI tools predict the next word in a sequence based on learned patterns in the data, therefore they are skilled at mimicking human language patterns. They lack a mechanism to verify the truthfulness of the content. If a model hasn’t seen specific information or lacks recent training data on a topic, it may fill in the gaps by generating an answer that sounds plausible based on similar topics. Also, if a prompt is too broad or involves complex topics, GenAI tools may mix together loosely related details in an attempt to answer the prompt, even if this leads to inaccuracies.
- **Data security and ownership:** anything users feed into a GenAI tool will typically be stored on a remote server. This is required for the operation of the service, as most users want to be able to refer back to previous interactions, and because companies want to be able to do quality control, troubleshooting and analysis. But stored data can also be analyzed for trends, contents or personal data and be used for further training of the GenAI model. In many cases, such data become the property of the owner of the GenAI tool, which has important implications for privacy, intellectual property and authorship rights.
- **Environmental effects:** the datacenters required to operate (generative) AI tools are resource intensive. The ecological footprint generated by the use of GenAI should be considered.
- **Working conditions:** Data labelling is crucial for training AI systems, meaning people assess images and texts to help the AI recognize different types of content and respond appropriately. These ‘trainers’ are often employed in low-wage countries and subject to some horrific content with limited psychological assistance for employees available.

This list is not exhaustive, though these are the most-cited drawbacks to GenAI tools. It is important for users to be aware that use of this technology comes with both positives and negatives. In the rest of this document, we will focus on practical implementation challenges in a higher education context, though at times we will refer back to some of the risks listed above.

It is important to always be aware of that GenAI tools are just that: tools. The tools in itself are neither good nor bad, they have some major advantages but also some major drawbacks and flaws. The way a user interacts with the tool and uses the outcome, determines the result and the tool's usefulness. This also means that the end-user is always responsible for the final product.

3. Challenges of GenAI in higher education

Higher education typically emphasizes knowledge acquisition, which presents significant challenges when using GenAI tools, especially in teaching approaches that focus heavily on written products (e.g., theses, programming assignments). This trend raises important questions about how students are assessed and what roles teachers and programme management should play³.

Each programme within the VU has defined a set of exit qualifications. In the course of a study programme, we ensure that students are taught the knowledge and skills to demonstrate their mastery of these exit qualifications. There is an extensive quality control system in place to ensure that all students who complete their study programme successfully earn the right to their degree and title. As developments in GenAI progress, it will become increasingly difficult to distinguish whether a student submits work that is based on their own thinking and work. When we consider [Bloom's taxonomy](#), on which our learning objectives are largely based, GenAI's strengths often lie in tasks related to remembering, understanding, applying, and analyzing, while its abilities in evaluating and creating are still evolving and may not match those of humans in all respects⁴. Yet, even if GenAI tools cannot convincingly perform all tasks yet, these tools will be developed further and will mimic human responses more and more accurately. As such, GenAI poses a direct threat to the validity of some of the assessment methods we use to ascertain whether students have mastered their study programmes' exit qualifications.

The traditional focus of academic education on results (grades) might give the impression to students that it does not matter how they arrive at the outcome. This opens the door to cutting corners. That in itself is not a new problem; however, it has become substantially easier to cut corners with easier access to GenAI tools. Keeping in mind the challenges GenAI poses to designing a valid assessment, one could argue that both **teaching and assessment methods need to focus on the process** (and not the outcome) of learning, more than ever. Course design should, therefore, emphasize the learning process (formative) rather than (summative) assessment alone. The VU vision on assessment⁵ also takes this approach by putting the focus on the learning process at the core of the vision. In part III and part IV of the BETA guidelines on use of GenAI, a selection of tips and examples is offered on how to implement more formative types of assessment.

3.1 Challenges of GenAI in assessment methods

Generative AI poses challenges to the validity of examinations in terms of **academic integrity** and **assessment accuracy**. These concerns can be relevant to any written assessment product, be it an essay, a report or programming code. The following list summarizes the main challenges GenAI poses to assessment but is far from exhaustive.

Academic integrity

- **Plagiarism and/or fraud:** In terms of plagiarism and/or fraud, it really does not matter whether a student cheated by copying other students' work, by letting another person do the work for them, or by copying the work of an inanimate object (GenAI). As long as the correct source is not provided and the student did not do the work themselves, this would be a clear case of fraud. The difficulty lies in determining whether or not the student actually used GenAI. Development on GenAI tools is progressing at a fast pace, and so called "AI humanizers" are being designed to make artificial intelligence systems produce outputs that are more relatable, empathetic and human-like.
- **Loss of critical thinking skills and problem-solving skills.** If a student uses GenAI to do their assignments and so pass their courses, they bypass learning processes aimed at teaching

³ Kohnke, L., Moorhouse, B. L., & Zou, D. (2023). ChatGPT for Language Teaching and Learning. *RELC Journal*, 54(2), 537-550. <https://doi.org/10.1177/00336882231162868> (Original work published 2023)

⁴ Esther Schagen (VU, 2023). Guidelines for Teachers of The Faculty of Social Sciences on Generative AI, via <https://assets-us-01.kc-usercontent.com/d8b6f1f5-816c-005b-1dc1-e363dd7ce9a5/c3fe9419-c736-4db6-93d9-217364021ee2/Guidelines%20for%20teachers%20of%20FSS%20on%20Generative%20AI.pdf>

⁵ VU assessment vision, 2022. <https://vu.nl/en/employee/teaching-at-vu-amsterdam/assessment-vision>

them how to function in their field, including fundamental knowledge and skills required to parse and value information and how to contribute to professional discussions and debates. Therefore, they forfeit the opportunities to learn skills that are invaluable in their future careers and which the academic world prides itself for. There is a need to redesign our teaching and assessment, to make sure we assess the right skills (and not whether a student knows how to use GenAI).

- **The input of sensitive data in (open) GenAI models**, for example personal information or confidential information regarding studies or organizational information collected in the context of internships, poses ethical concerns regarding data privacy.

When designing assessments, please be aware of the following limitations GenAI tools have:

- **They cannot produce originality.** GenAI models are trained using large datasets and use this as the basis for their output. That way, GenAI models are capable of producing text that is unique in appearance without being original, as it will always be a result of combining existing data. This leads to a significant misconception: that AI-generated outputs are innovative or original in the same way that human-created works can be.
- **Very specific and contextual information may be beyond the abilities of a GenAI tool.** GenAI tools are trained using large datasets. Therefore, if an assignment draws on students' personal experiences or is built around a very specific target group, GenAI might not be able to help a student out, given the lack of relevant material in the tool's training data.
- **The quality of GenAI-generated responses may vary significantly.** Simply using a prompt and copying the outcome may either lead to acceptable (or even outstanding) work, but might also lead to subpar results if the student does not thoroughly check the output.

3.2 Challenges of GenAI detection software

With the breakthroughs in GenAI, there has also been a countermovement developing **AI detection software**. These tools can often determine correctly that GenAI has been used but also give rise to false positives. Since detection software is based on statistics, just as generative AI itself, there will never be a way of stating with 100% certainty that text is either generated by AI, or not. Given this weakness, use of GenAI detection software as evidence of fraud is strongly discouraged and is not considered proof in fraud cases. It may, however, be a starting point for a conversation with the student about their suspected GenAI use.

3.3 The distinction between Assessment of Learning and Assessment for/as Learning

There is a distinction between Assessment *of* Learning and Assessment *for/as* Learning⁶. Assessment of Learning is mainly summative, assessing whether a student meets final exit qualifications. Assessment for/as Learning is mainly formative, with assessment and feedback contributing to the learning process. Both are necessary to get a good picture of the students' learning process and its outcomes. For the design of assessment and the use of GenAI, this distinction is important⁷

Assessment of Learning produces a grade or pass/fail and assesses whether a student is meeting the learning objectives. As a teacher, this is where you might want to impose restrictions (control) on the use of GenAI to ensure that learning objectives are met. To create a valid assessment, both the assessment form and the assessment environment are important. For example, if particular writing skills (e.g. phrasing scientifically versus colloquial) are part of the learning objectives, the use of GenAI may not be allowed in the assessment of these particular skills. If writing skills are not part of the learning objectives, it is wise to choose an alternative assessment format that is not based on writing (e.g. presentation, discussion, debate, pitch), where use of GenAI does not risk the validity of

⁶ Black & William, 1999

⁷ Liu, D., & Bridgemen, A. (2023, July 12). What to do about assessments if we can't out-design or out-run AI? – Teaching@Sydney

the assessment. If it is necessary to ensure that a student does not use external sources, such as use of GenAI, a controlled environment is key. It is also possible to imagine situations where a semi-controlled environment is possible, where the focus of assessment is on the process and transparency around students' use of GenAI.

Assessment for/as Learning focuses on the learning process and provides feedback that students can use to adjust their learning to increase their study success. Thereby, assessment for/as learning lends itself perfectly to learning how to use GenAI in the right way (AI Literacy), as it prompts reflection on how students accomplish results. GenAI thus becomes a thinking and working tool of students and future professionals⁸. In this process, through guidance, learning to prompt, critical analysis and discussion, and attention to academic integrity, lecturers and fellow students maximize the stimulation of student learning through assignments. Assessment for/as Learning takes place during classes, group-meetings and at home. Given that the learning environment is partly informal and unsupervised, it is not desirable to prohibit the use of GenAI. The use of GenAI is the student's own responsibility.

By determining the type of assessment, examiners can identify the assessment environment and the acceptable use of GenAI during the assessment. In the BETA policy for use of GenAI, we will introduce a classification for the use of GenAI for each type of assessment.

4. What do students need to learn about GenAI?

Given the role GenAI has already assumed in society, our next logical step is to establish how to integrate the potential uses and risks of GenAI into our teaching. Part of our job is to prepare students for their future, be it a career in academia or in the broader society. Wherever a student ends up, they'll likely encounter some form of GenAI. Therefore, one might argue that equipping students with **the skills to use GenAI tools adequately, correctly and responsibly** is part of the programmes' job (we call this **AI literacy**). Of course, some argue that this is a short-term issue, since new generations will grow up with GenAI as part of their everyday life, just like the current generation is growing up with other digital tools. However, consider that a large part of the current younger generation does not know how to **distinguish between fact and fake news**, how to **correctly check sources**, and how to **critically assess evidence presented**. As we have reasoned before when listing the pitfalls of GenAI, these skills will become of even greater importance with the further rise of GenAI. We cannot rely on students being taught this in elementary and secondary school. The next skill to teach our students in relation to GenAI is ethics.

While **ethics and academic integrity** have always been objectives in academic study programmes, the challenge increases with the use of GenAI. One of the issues regarding academic integrity is the copy/pasting of generated output and passing it on as someone's own work. This would be considered fraud by the rules of the examination board, as it makes it impossible to judge the student's own knowledge and skills. Also, since GenAI typically relies on a learning model, where data that is fed into the system can be saved for future training data sets, **personal or sensitive data** must not be entered into GenAI tools in any circumstance. However, research suggests that as much as 55% of GenAI inputs contains such data⁹. This underlines the need to teach our students what constitutes sensitive information, and why it is so important never to enter it into GenAI tools. The high percentage of sensitive GenAI input suggests that our society does not yet treat sensitive information with appropriate care and restraint. Additionally, the definition of sensitive data might not be consistently formulated or accepted in the same way across society. From preliminary inventories of some Science students, it seems that many students use ChatGPT as a substitute for Google; they rely on it for factual information, as an introduction to certain topics that they find difficult. Without further fact-checking, this poses a risk to academic integrity.

⁸ <https://vu.nl/nl/medewerker/didactiek/ai-bestendig-toetsen-met-de-tweebaansbenadering>

⁹ <https://resources.menlosecurity.com/reports/the-continued-impact-of-generative-ai-on-security-posture>

Microsoft Copilot

To eliminate problems with private and sensitive data, VU has a license agreement for the use of Microsoft Copilot, for anyone with a VUID. Data are protected within VU, meaning data will not be used for training the AI model. Microsoft Copilot can also save a prompt/chat history, which students can download as a log (proof of how they used GenAI in their assignment). Teachers can ask students to use this specific tool, they cannot ask or require the use of other tools (such as ChatGPT) for which VU has no license agreement. The license also helps in avoiding any student inequalities, since all students have equal access to this tool. However, keep in mind that there are other GenAI tools, where students who have the means to purchase the paid version or elect to share private information might still be put at an advantage over other students.

Both students and staff need to be aware of the possible consequences of entering and storing data. More explicit: the faculty board of BETA VU states that no personal or sensitive data may be entered into GenAI tools in any circumstance, including Microsoft Copilot. Even though the data entered into Copilot are kept within the VU container, data is still stored. This makes the data susceptible to leaks through hacks, for example. Once data is stolen it can be sold and linked to larger data sets. This can lead to seemingly innocuous data linking back to individuals. When sensitive data have been entered into Copilot as a 'safe' tool, the risks are even greater.

Since the use of GenAI in education will be scrutinized by examination boards we need to make abundantly clear to students **what does and does not constitute fraud**. VU BETA qualifies fraud as: *“any act or omission by a student that partially or entirely precludes an accurate assessment of their knowledge, understanding and skills, or those of another student.”* (See article 19 of the Rules and Guidelines of the examination board BETA VU¹⁰ for further explanations of the terms 'fraud' and 'plagiarism'.) Evidence seems to suggest that there appears to be limited awareness as to what constitutes fraud amongst students¹¹. There is a big role for institutions in preventing fraud, by better guiding students and by explicating what procedures are when fraud is suspected and/or established. A higher degree of awareness among the student community and more efficient information campaigns seem to be the key to limiting fraud in education.

4.1 Generative AI as a writing aid

Another way in which students are using GenAI tools is as writing aides; asking GenAI tools to improve their own text. This is similar to more traditional computer assisted writing aids that have been in use for years, for example in Microsoft Word, helping them improve grammar and spelling, but also sentence structure and offering other forms of text enhancement. While these tools rely on pattern recognition for scripted suggestions for text improvement, and not creating new content per se, some of these tools are now offering GenAI functionalities (sometimes in a paid version only). Therefore, we feel it is right to include a policy on these types of tools as well. It is almost impossible to banish tools for revision, since it is almost impossible to prove the (ab)use of a writing aid. Therefore, the better option is to guide students to use writing aids responsibly, whether they are GenAI or simply computer assisted writing aids (Microsoft Word, Grammarly, DeepL Write). The conditions under which students may use GenAI as a writing aid are stated in Part II: BETA policy on the use of GenAI. Use of these tools may be particularly beneficial for non-native speakers of the programme language, or those who struggle with language fluency and/or dyslexia, thus reducing student inequalities.

¹⁰ Rules and Guidelines examination board BETA VU 2024-2025 <https://assets-us-01.kc-usercontent.com/d8b6f1f5-816c-005b-1dc1-e363dd7ce9a5/0dcf44f3-3847-4ad8-be63-30271239c01b/Rules%20and%20Guidelines%20Examination%20Board%202024-2025.pdf>

¹¹https://www.cimea.it/Upload/Documenti/Fraud5_Student_awareness_on_fraud_in_education_singole.pdf