Working with GenAI within FGGA

Version 2.0

1 July 2026

Generative Artificial Intelligence (GenAI) is being increasingly adopted within FGGA. The technology itself is also developing. GenAI can potentially improve the quality and efficiency of our faculty, but we must be ready to work with it deliberatively and with awareness.

Within FGGA, we want to take a proactive and positive approach to GenAI. This document provides practical guidance for the use of this technology in our faculty in the form of a Stoplight Framework. The guidance of the Stoplight is based on legal and technical grounds, but it is only the implementation of overarching principles which guide our approach.

Our approach for the use of GenAI within FGGA is that **we train, inform and support our employees so that:**

1. **We entrust our employees** – We seek to support our staff to make *informed decisions.* Our staff are capable of responsibly exploring GenAI tools with the right support. Thus, we will focus on training, sharing information, and providing safe working environments.
2. **We explore the opportunities** – Responsible experimentation helps us better understand both the risks and the potential gains of GenAI. Therefore we will stimulate pilots and encourage experimentation. We want to know what this technology *can* do. To ensure we learn from these experiments, we will focus on capturing and sharing lessons learned.
3. **We work transparently and align with public values** – We prioritize transparency in how and where GenAI is used. We respect privacy regulations and intellectual integrity in all GenAI use.
4. **We remain aware –** We recognize that GenAI also brings real risks and limitations. Therefore, we will remain critical, evaluate tools carefully, and set clear boundaries to ensure ethical and purposeful use.

To properly entrust our staff, we need to provide them with sufficient support and guidance. We need a **programmatic approach** at FGGA that includes pilots, training, and clear information which is consistent. We also need to give explicit permission for staff to explore GenAI tools under defined and safe conditions.

Changes since v1.0 June 2025

Preface to the updated stoplight framework

The further adoption of Generative AI at Leiden University is clear. The ISSC reports rapidly increasing frequency of messages being sent back and forth to commercial AI providers: in December 2025, there were more than 600,000 weekly messages exchanged with commercial AI providers, but by March 2026, this number had grown to more than 1,000,000.

Given that the University has no signed agreements with these commerical parties, we have simply no management of the data, privacy, and legal risks entailed.

The original Stoplight framework was intended better clarify, in simple terms, under what conditions it would be acceptable to work with AI tools. The tagline of this policy – *what is not forbidden is allowed* – also means that there are many legitimate and safe ways to interact with commercial parties with whom the university has no processing agreements.

However, without proper IT infrastructure, the vast majority of the daily work of our faculty is not suitable to be shared with external parties. For example, management-level strategic information, information generated on personnel emails which contain personal information, financial records, reports, and analysis, etc., would all be considered in the red category of the original stop-right framework. Publishing or sharing this data with a commercial AI provider is equivalent to handing over ownership of this information to an unregulated foreign commercial party.

Recently, FGGA began a pilot to create its own digital shell to safely and securely process data with a much higher degree of digital sovereignty, deploy a variety of market-standard AI models to overcome this limitation. This tool is called LUChat. As a result, the Stoplight now distinguishes between LUChat and all other commerical tools. Should the University conclude agreements for the processing of data with commercial party, this Framework would need to be edited again.

The two other major sections of the Framework (Understand the technology, and; Understand the impact) are unchanged. While GenAI is evolving, there are no material changes to the considerations posed. Meanwhile, the impact of GenAI can only be named in such a Framework.

FGGA GenAI Framework

Understand the data, understand the technology, understand the impact. What is not forbidden is allowed.

Below is a summary of the FGGA framework. It poses 3 questions you should answer before using AI.

🟢 If the answers to all these questions are positive, *go ahead!* You have quite a lot of freedom in the tools you can use for these use cases.

🟠 If you answer 'unknown' or 'undecided', you must *only use LUChat* and you may reach out for a discussion with us *responsibledata@fgga.leidenuniv.nl*

🔴 If the answer to any of these questions is negative, *stop, and see if you can mitigate the risks, for example, by anonymising or pseudononymising the data.*

Understand the data		**Understand the technology**	**Understand the impact**	**What does that mean for me?**	
Is it permissible to process this data with a GenAI?		**Will GenAI deliver good results?**	**Am I comfortable with the costs/side effects?**	**Which Tools may I use?**	**What are my responsibilities as a user?**
External GenAI tools	**LUChat**				
🟢 Public, CC00/CC-BY, non-personal, non-confidential data	🟢 Public, open documents	🟢 Content creation where review is easy, accuracy is not essential, and novelty is not important.	🟢 The value outweighs the cost and potential collateral impact.	You may use any tool here. However, default to LUChat or Microsoft Copilot.	You must uphold copyright. You must accept responsibility for the generated content you use. You should inform the recipient if your work is **not substantially altered** from an AI generated output.
🟠 Internal work documents	🟢 (All) Internal work documents	🟠 Not relying on GenAI for true novelty or perfect accuracy.	🟠 The effort of checking GenAI's work is lower than doing the work myself.	You may only use Microsoft Copilot or LUChat .	Same as above.
🔴 Personal data; Special category data (GDPR); autonomous decision making purposes.	🟢 Sensitive and special category data.	🔴 (Autonomous) high-stakes financial, legal, strategic decision; producing replicable scientific work.	🔴 I feel personally uncomfortable with the environmental, social or ethical impact.	You may only use LUChat, within the purposes allowed by the AI Act.	Check carefully whether the purposes of your use are allowed by the AI Act. Otherwise, you must not use GenAI.

Section 1: Understanding the Data

Data has to clear 3 hurdles before it can be shared with a GenAI.

1. Is it *permitted by regulation* to share?

> 🔴 All Tools
>
> The following purposes are explicitly forbidden by the AI Act. The AI Act forbids the use of AI technology for these purposes in general, regardless of how the data is handled or processed.

In all cases, the AI Act and GDPR apply to our use of digital technology. These regulations strictly forbid certain practices. This holds true even if a certain tool may be permissible to use. There are prohibited practices under both GDPR and the AI Act:

Category	**Definition**	**Within FGGA, hypothetical example:**
Processing Special Categories of Personal Data without Consent	Includes data revealing racial or ethnic origin, political opinions, religious beliefs, or health/sexual orientation, unless explicit consent is given.	- Without consent, processing transcripts from interviews and asking for a summary of common themes. - Scraping a social media profile to attempt and place the user on the political spectrum
Automated Decision-Making without Human Involvement	Decisions made solely by algorithms that significantly affect individuals are restricted unless consented to or legally authorized.	- Feeding CVs into a GenAI and automatically accepting the recommendation on which candidate to hire. - Automatically grading a paper without any review or independent human thinking.
Manipulating Behaviour	Use of AI systems to manipulate human behaviour subliminally or subvert or impair autonomy.	- Using automatic AI-generated communications to manipulate students into following study routes.
Monitoring and/or manipulating emotions	Use of AI students to infer emotions in a workplace or educational setting.	- Using AI to infer students' emotional state (for example, during an exam)
Social Scoring by Governments	AI used by public authorities for social scoring that results in unjust or harmful treatment.	Not strictly applicable to a University, but it might still be foreseeable that we become complicit in this activity if we uncritically share data.
Real-Time Remote Biometric Identification	Real-time facial recognition or other biometric identification in public spaces is generally banned	- Automatic detection for staff/students for admittance into LU buildings.

	(exceptions for law enforcement may apply).	

2. Is it *permitted by patent and copyright law* to share?[1]

LUChat
All External Commerical Tools

Sharing this type of information with an external tool effectively transfers ownership of the content to the tool provider. However, LUChat processes information statelessly within cloud servers to which only LU has access. As such, the data is never shared or transferred.

Does the data set you are working with (the prompt, attachments, or links) contain confidential or secret business information? For example, does it contain financial forecasts, contract drafts, procurement data, or information which is protected by any kind of confidentially agreements?

Is the content protected by a copyright licence?[2] As a general rule of thumb, if you created the content you have the right to publish it under Creative Commons lisence and thereafter share it with GenAI:

- Published academic works (such as academic papers from a journal) bear restritive copyright lisences which make it forbidden to feed into a GenAI tool;
- However, you retain the copyright of your own published academic work;[3]
- Teaching materials (slides, lecture notes, etc.) remain the copyright property of the University, but the University gives you the right to share these under a Creative Commons lisence;[4]
- As reflected in the FGGA Data Protocol, Leiden University also explicitly entitles researchers to share their research data under the same licences;[5]

[1] See also the University Library's Copyright Information Office https://www.library.universiteitleiden.nl/researchers/copyright-information-office 25 June, 2025
[2] See the full copyright policy of the University (updated in 2025) https://www.organisatiegids.universiteitleiden.nl/binaries/content/assets/algemeen/reglementen/leiden-university-employer-copyright-regulations_4.2.2025.pdf June 25, 2025
[3] https://www.library.universiteitleiden.nl/researchers/copyright-information-office/copyright-on-teaching-materials 25 June, 2025
[4] See this explainer for the different types of Creative Commons lisences. https://creativecommons.org/share-your-work/cclicenses/ 25 June, 2025
[5] https://www.library.universiteitleiden.nl/researchers/open-access/background-and-policies 25 June, 2025

3. Is it *wise* to share?

All Tools

Just because AI *can* do something, does not mean it should always be entrusted to do it. Regardless of the tool, AI can introduce errors, strip important nuance, and errode quality in ways are difficult to spot but which might have a significant impact on others.

Even when sharing is legal and permissible, it may still be unwise. Ask yourself the following:

- **Could the output be misunderstood or misused?** If context is lost, the GenAI could generate outputs that are misleading.
- **Will the output be relied on by others?** If your output will be published, cited, or used for decision-making, apply extra caution. Review and validate the results thoroughly.
- **Is there a reputational risk to yourself, colleagues, the university?** Consider how sharing this content with an AI tool — or using AI-generated output — might reflect on you or your institution if made public.

When in doubt, seek a second opinion or choose for more human oversight.

4. There is lots of information which is *perfectly OK* to share with GenAI

LUChat
All Tools

LUChat does not store, transmit, reproduce or retain the information you share with it. As a result, it avoids the increasingly restrive application of copyright to publicly available information.

For instance, information that is already publicly available — such as reports, interviews, blog posts, or datasets published with Creative Commons licences — can generally be used in all AI tools with fewer concerns about intellectual property, confidentiality or data protection. This is, however, a general rule of thumb, and the landscape is changing.

However, as with any form of reuse, it remains your responsibility to ensure that such material is used ethically and in compliance with relevant licenses. It is always your responsibility to properly reference your sources.

Section 2: Understanding the Technology

Generative AI works by **predicting** an expected outcome based on historical patterns. In comparison, a digital calculator follows strict, predefined rules to deliver precise and replicable answers based on mathematical logic. That means that this GenAI technology is ultimately probability-based rather than rule-based.

This makes GenAI **good at** tasks that involve working with natural language, like summarising and rewriting, and at working with well defined formats (like emails, presentations, papers, lists) where there are historical examples. It also makes GenAI good at generative visual output where there is some room for error (such pictures, emojis, graphs).

At the same time, the probabilistic nature of GenAI means it is **less effective** for tasks that demand precise, replicable and complex interpretative operations. GenAI is still 'guessing' the outcome of mathematical operations and thus it is an unreliable assistant on data analysis. GenAI is also trained on historical data, meaning it still struggles to generate truly new or creative insights that it has never seen before.

It is therefore important to match the type of task to GenAI's strengths, and always apply human oversight to the results.

GenAI is Good at:	GenAI is OK at:	GenAI is Poor at:
· Summarising existing information, such as writing summaries of public documents, meetings, papers and webpages · Rewriting text for clarity, either to expand or contract · Generating creative ideas · Formatting emails or reports	· Drafting routine policy text (non-binding) · Generically analyzing data sets and documents · Translating general content (non-sensitive) · Answering general-purpose knowledge questions · Generating slides for internal presentations	· Creating genuinely new ideas · Making legal or financial decisions · Interpreting nuanced policy implications · Performing accurate data analysis · Producing the replicable results

Section 3: Understanding the Broader Impact of GenAI

Beyond the clear and exciting benefits that GenAI brings, there are also clear costs, thought these may be indirect. Even if how you want to use AI is safe, legal and appropriate, it may still be undesirable in certain contexts because of the indirect impact this technology has on society and the environment.

1. Environmental Impact

Using large-scale AI systems consumes significant amounts of energy and water.[6] [7] Each interaction may seem low-cost, but widespread use across an institution can lead to substantial ecological footprints—especially if tools are hosted in regions with limited renewable infrastructure or scarce water resources.

2. Intellectual Property and Institutional Sovereignty

Leiden University generates a lot of high-quality data. When we share that data with third-party AI tools, we are typically contributing to the training data or internal knowledge of commercial providers. This raises the quality of *their* product but does not necessarily provide a direct benefit to us. This raises concerns about:

- **Knowledge leakage** to private firms or non-EU jurisdictions. Information shared with GenAI tools almost always becomes proprietary knowledge of the developers.
- **Vendor lock-in**, where our most important knowledge and workflows become dependent on external (often opaque and expensive) commercial tools.
- **Loss of control** over how academic or strategic content is used and commercialised.

3. Ethical Concerns

Because of GenAI's *averaging* and *middling* method of operation, it is feasible that over-reliance on GenAI lowers the overall quality of work and skill that human agents develop.

- **De-skilling and disintermediation**: Replacing human work with AI-generated outputs may limit learning opportunities for students and remove expert judgment from decision processes.
- **Bias amplification**: AI models reflect historical biases in their training data, which can lead to the reinforcement of stereotypes or marginalisation of underrepresented perspectives.
- **Homogenisation of thought**: Because AI tends to reproduce the most statistically probable answer, it risks eroding intellectual ingenuity and discouraging original or dissenting viewpoints.

[6] https://e360.yale.edu/features/artificial-intelligence-climate-energy-emissions
[7] https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about

Practical Examples of When to Use AI

- You may use generative AI.
- There are restrictions on which generative AI tool you may use.
- It is forbidden to process this information (at all or using GenAI).

GenAI Use Cases: Permissibility vs. Concerns

Below are some examples of potential use cases in FGGA:

Use Case	With Commercial AI Tools	With LUChat	Remain Aware of Potential Risks / Concerns
Summarizing a public policy document, from the University or a government Ministry	Publicly available, non-sensitive	Publicly available, non-sensitive	May oversimplify or misrepresent key points if not reviewed by a subject expert
Expand provided bullet points into a formal academic prose, for example for writing a memo to the Faculty Board	Permitted **if** work is own IP and it contains no personal or confidential information	Permitted	May introduce tone mismatch or misleading phrasing, harming your credibility
Teacher works with GenAI to write a course syllabus	Permitted **if** no personal data	Permitted	GenAI may alter tone or clarity, or remove pedagogical intent
Creating social media captions based on a press release	Permitted **if no** personal data in original text	Permitted	May introduce tone mismatch or misleading phrasing, harming institutional credibility
Drafting FAQs and explainer emails from internal policies	Permitted **if** policies are not confidential and contain no personal information	Permitted	GenAI may generalize or overstate responses, leading to miscommunication
Sharing lecture notes and slides, and generating example questions for an exam or quiz	Permitted **if** non-sensitive information, no student information, and content is published under CC license	Permitted	Suggestions may be inappropriate, inaccurate, or poorly aligned with learning goals
Drafting an internal memo from meeting notes	Permitted **if** no personal or	Permitted	Possible distortion of intent, structure, or strategic messaging

	confidential information		
Rewriting job vacancy text	⚠🟢 Public, non-sensitive information	🟢 Permitted	May introduce bias or reduce specificity, harming inclusivity or clarity
Rewriting a grant proposal	⚠🟢 Permissible **if** it contains no confidential or third-party content	🟢 Permitted	May misrepresent finer points of research; may lower the quality of the prose
Creating meeting notes from a transcription	🟠 Permissible **if** anonymized **or** explicit consent has been documented	🟢 Permitted	Possible misattribution, distortion of intent, misquoting. May omit or distort key details
Summarizing student evaluations by course	🟠 Permissible **if** anonymized and aggregated (most are by default)	🟢 Permitted	Risk of de-anonymization or misrepresenting perspectives
Suggesting edits to a GROW meeting summary	⚠🟠 Permissible **if** anonymized and non-sensitive	🟢 Permitted	Risk of hollow or generic suggestions that undermine authenticity
Summarizing interview transcripts conducted for research	⚠🟠 Strict preconditions for consent must first be met. Contains identifiable or sensitive data	🟢 Permitted	Forbidden unless explicit consent is documented, and data is fully anonymized
Drafting a disciplinary letter for a student or staff member	🔴 Contains personal, sensitive or legal information	⚠🟢 Permitted **if** the use of GenAI is documented in process manuals	Forbidden under GDPR; risk of legal liability and reputational harm
Select a winning candidate from provided CVs or letters of intent	🔴 CVs contain personal information. Involves profiling and automated decision-making	🔴 Forbidden under the AI Act	Forbidden under the AI Act unless exceptional pre-conditions and human oversight are met
Transcribing meetings regarding health-related disclosures	🔴 Special category personal data under GDPR	⚠🟢 Permitted **if** informed consent is documented	Forbidden unless explicit consent is obtained and additional safeguards are in place
Summarizing meeting minutes from confidential labor negotiations	🔴 Strategic and sensitive institutional content	⚠🟢 Permitted **if** informed consent is documented	Forbidden due to confidentiality and legal obligations
Generating a *definitive* grade or assessment based on a written assignment	⚠🔴 Without sufficient human oversight breaches AI Act	⚠🔴 Without sufficient human oversight breaches AI Act	Forbidden if used without substantive human judgment or oversight

Teacher uploads student papers for feedback or evaluation	🔴 Work remains intellectual property of student	🟢 Permitted	Forbidden without informed consent, opt-out and registration. LU takes an ethical stance here