Recommendations for the Use of Artificial Intelligence Tools at UP

A Practical Guide for Teaching, Science and Research, Administration, and Support Activities

Palacký University Olomouc

Version of 4 May 2026

How to use this document

This document is intended as practical guidance for UP employees and students: it helps them quickly assess the purpose of AI use, the type of data entered, the operating mode of the tool, and the degree of human oversight. It is not a catalogue of all existing services or a manual for a specific commercial product.

If you need to make a quick decision, start with these sections:

1. Essential guidance for every UP employee and student.
2. The data traffic light: green, yellow, and red zones.
3. Permitted, conditional, and prohibited uses of AI.
4. Practical scenarios for teaching, research, and administration.
5. What to do in the event of an incident or doubt.

> **One sentence to remember**
> Do not enter anything into an AI tool whose data-processing rules you cannot verify (all online tools) that you would not publish on the web without hesitation.

Preamble

This document provides a practical framework for the safe, responsible, and controllable use of artificial intelligence at Palacký University Olomouc. It is a recommended minimum framework for the whole of UP; faculties and other units may specify it further or make it stricter according to their needs.

The document builds on previous UP recommendations on the use of AI, publicly available UP documents on personal data protection, ethics, and research integrity, related university regulations, and the European framework for the use of artificial intelligence, especially Regulation (EU) 2024/1689 of the European Parliament and of the Council, known as the AI Act.

Binding effect in specific situations arises mainly through internal regulations, syllabi, assignment instructions, publication rules, grant conditions, contractual obligations, and decisions of the relevant units. The current version of this recommendation, the FAQ, and related methodological materials will be published on the AI at UP web portal.

The document is technology- and media-neutral. It applies to chatbots, translators, image, audio, and video generators, data-analysis tools, office-software assistants, programming assistants, search assistants, AI agents, AI systems embedded in information systems, local models, API services, and specialised research or administrative tools.

Essential guidance for every UP employee and student

1. A human being is responsible for the output. AI may be a support tool, not an author, decision-maker, or substitute for expert judgement.

2. Assess the data first. Before using AI, determine whether you are working with public, internal, personal, sensitive, contractually protected, or research data.

3. Internal and non-public data belong only in a permitted mode. It is not enough that a tool is available, popular, or built into an application. The decisive factors are its contractual, security, and operational conditions.

4. Do not enter red-zone data into public tools. This includes, in particular, sensitive personal data, health data, payroll data, non-public examination materials, non-public contracts, unpublished research data, passwords, API keys, and access tokens.

5. Always verify AI outputs. AI may hallucinate, omit essential information, reproduce bias, create false citations, or formulate untruths convincingly.

6. Disclose the use of AI. Follow the rules of the course, faculty, publisher, grant agency, or workplace.

7. Synthetic images, voices, and videos require special caution. Creating deepfakes, false recordings, unsolicited likenesses, or sexually explicit synthetic materials is unacceptable at UP.

8. If an error occurs, stop the work. If you have accidentally entered sensitive or non-public data into AI, stop processing, note the basic circumstances, and contact your supervisor or the responsible person at your workplace. Details are provided in Sections 4.3 and 12.

1. Basic principles

1.1 AI is a support tool, not a substitute for responsibility

AI can help with suggestions, structuring texts, language editing, translation, summarisation, creating working drafts, programming, data analysis, creating visuals, preparing teaching materials, preliminary literature searches, or administrative support. However, it should not make the final decision in situations that have legal, study-related, employment-law, ethical, security, health, or scientific impact.

The user of AI is responsible in particular for:

- the suitability of the selected tool and mode of use,
- the legitimacy of entering the data,
- the factual correctness of the output,
- disclosing the use of AI where required,
- complying with rules on personal data protection, copyright, confidentiality, academic integrity, and cybersecurity.

1.2 AI must not be listed as an author

Generative AI has neither authorship nor responsibility for content. In academic, publishing, and administrative practice it is therefore not listed as an author or co-author. It may be listed as a tool used, an aid, or part of the methodology, if required by the rules of the publisher, grant agency, faculty, course, or workplace.

For scientific texts, it is advisable to describe in particular:

- which tool or model was used,
- for what purpose it was used,
- whether it affected only language editing or also the substantive structure, analysis, code, visualisation, or interpretation,
- how human review of the output was carried out.

1.3 The purpose of use, the data, and the tool mode are decisive

The same tool may be safe for work with a public text and at the same time unacceptable for work with personal data, a non-public manuscript, an examination assignment, or an internal contract. This document therefore addresses not only the question “Which tool may I use?”, but also the question “Under what conditions and with what data may I use the tool?”

1.4 Transparency and academic integrity

The use of AI must be appropriately disclosed where AI has influenced the content, method, explanation, analysis, or form of the output. Undisclosed deliberate use of AI may constitute an ethical breach, a breach of course rules, a breach of publication rules, or a breach of research-integrity rules.

Transparency does not mean reporting every trivial typo corrected by an automated tool. It does mean disclosing significant AI assistance that affected the content, structure, argumentation, translation, calculation, visualisation, code, research procedure, or interpretation.

1.5 AI literacy is a minimum condition for responsible use

Everyone who uses AI in work or study at UP should understand at least the basic properties and limits of the tool. Minimum AI literacy includes the ability to:

- recognise that AI may produce inaccurate, fabricated, biased, or incomplete outputs,
- assess the sensitivity of the data entered,
- distinguish between public, internal, personal, and highly sensitive use,
- know when human review or consultation is necessary,
- recognise risks associated with synthetic media, hallucinations, bias, prompt injection, data leakage, and automation bias,
- distinguish between a tool for initial drafts, inspiration, and auxiliary text processing and a tool suitable for decision-making or scientific output.

2. Assess the data first: a traffic-light system for working with AI

Before every use of AI, ask yourself: What exactly am I entering into the tool, and why? If the objective can be achieved without entering sensitive or non-public data, choose a safer option: anonymisation, a synthetic example, local processing, an internal tool, an aggregated description, or consultation.

Zone	Type of data	Where the data may be entered	Typical examples
Green	Public information, your own ideas, and working drafts without personal data, internal information, or non-public know-how.	Into ordinary public tools, provided you do not breach copyright, contractual confidentiality, or workplace rules.	A public article, a general outline, a public abstract, a general e-mail without personal data, language editing of a public text.
Yellow	Internal or non-public working materials without high sensitivity; pseudonymised or aggregated materials; materials with limited impact if disclosed.	Only into a mode permitted by the university, contractually covered, or otherwise securely configured.	An internal methodological text, a working presentation, a pseudonymised dataset, anonymised minutes, an internal draft without sensitive data.
Red	Sensitive personal data, non-public examination materials, non-public contracts, trade secrets, unpublished research data, access credentials, and security secrets.	Do not enter into public AI tools. Handle locally, in an approved internal environment, or after consultation with the Vice-Rector for Research, Creative Activities and Knowledge Transfer, who will involve the DPO, Legal Department, or Computer Centre according to the nature of the matter.	Czech personal identification numbers (rodná čísla), health data, payroll data, student lists, examination results, passwords, API keys, access tokens, patentable ideas, manuscripts under review, data from ongoing research.

Important clarification: The fact that a text contains no personal data does not automatically mean that it may be entered into a public tool. It may contain non-public know-how, contractually protected information, copyrighted material, a non-public research output, or reputationally sensitive content.

3. Permitted, conditional, and prohibited uses of AI

3.1 Permitted mode of use

A permitted mode of use means a specific combination of tool, purpose, data type, contractual terms, security configuration, and responsibilities. This document therefore does not contain a static list of AI tools: what is safe or unacceptable is not the name of the service itself, but its specific use in a specific mode.

A permitted mode of use has the following characteristics:

- it is intended for a specific type of work and data,
- the contractual terms and the provider's responsibilities are clarified,
- it is clear whether inputs and outputs are used for training or improving the model,
- data retention and logging are described,
- access control and appropriate security elements are in place,
- for personal data, the GDPR regime is clarified,
- it is stated who approves, manages, and updates it.

For public chatbots and APIs without contractual and security guarantees, assume that they are not intended for the yellow or red zones.

3.2 Practical checklist for the yellow zone

Before using AI for internal or non-public data, at least the following must be clear:

- in the default settings, the data are not used for further model training, or this option has been contractually and technically disabled (opt-out),
- the retention period for inputs, outputs, and logs is known,
- it is clear where the data are stored and who can access them,
- the tool allows access and roles to be managed,
- it is clarified whether a data-processing agreement is required,
- it is clarified whether the use requires a data protection impact assessment,
- there is a procedure for reporting incidents,
- users have been instructed on how to work with the tool.

Opt-out settings in a personal or free account, or the mere availability of an AI function in e-mail, a browser, or office software, do not by themselves create a permitted mode for internal or personal data.

3.3 APIs, keys, tokens, and automation

Additional risks arise when AI is used via an API. It is necessary to distinguish a token as a unit of billing or text length from an API token / API key / access token as a secret access credential, which belongs to the red zone.

Basic rules:

- Never enter API keys, access tokens, or secret configuration values into a public AI tool.
- Do not store API keys in source code, public repositories, shared documents, or prompts.
- For teaching, use test accounts and demonstrations without real internal data.
- For research and administration, use centrally approved accounts, access control, and auditable settings.
- For automation and AI agents, set limits, human approval, and logging, especially if the tool can send e-mails, edit documents, write to systems, run scripts, or access data.
- When using an API, it must be clear whether the provider stores inputs and outputs, for how long it stores them, and whether it uses them for training or improving the model.

3.4 AI agents and connected tools

An AI agent is a system that not only generates an answer, but may plan steps, use external tools, read files, search, run scripts, work with a calendar, e-mail, or a database. The more permissions the agent has, the higher the risk.

For AI agents:

- grant only the minimum necessary permissions,
- separate testing environments from production environments,
- do not allow an agent to perform irreversible actions without oversight,
- keep logs of steps and outputs,
- do not connect an agent to non-public data without an approved mode,
- allow for the risk of prompt injection, where malicious text in a document, web page, or e-mail instructs the tool to behave in an unintended way.

4. Personal data, GDPR, and incidents

4.1 Basic rules for personal data

When working with AI, the general rules for personal data protection must be followed. In typical situations, UP acts as the personal data controller. Each user is responsible for entering into AI only data that they are authorised to use, and only to the extent necessary for the given purpose.

Practical rules:

- enter only the necessary minimum of data,
- prefer anonymisation or aggregation,
- do not use public tools for sensitive personal data,
- for personal data, always assess whether the mode is permitted,
- for new or more extensive processing, consult the DPO,
- do not use AI for decisions about persons without legal and methodological assessment.

4.2 Examples of incidents

An incident may include in particular:

- entering personal, study-related, health, payroll, or other sensitive data into a public AI tool,
- entering non-public contractual, research, publication, or examination materials into a service without contractual protection,
- entering a password, API key, access token, or other secret value into a prompt,
- creating or disseminating a synthetic image, voice, or video of a person without a legitimate purpose and consent,
- using an AI output or AI detector as a basis for a decision without verification and without an appropriate process.

4.3 What to do in the event of an incident

If you accidentally enter inappropriate data into AI or have doubts, proceed as follows:

1. Stop further processing. Do not continue the conversation and do not enter any further data.

2. Preserve basic information. Note the time, tool, type of data entered, whether the use was public or in an approved mode, and who had access to the output.

3. Do not disseminate the output. Do not download or send the generated content to other persons.

4. Contact your supervisor or the responsible person at your workplace and the Vice-Rector for Research, Creative Activities and Knowledge Transfer. The Vice-Rector will perform triage and, depending on the nature of the matter, involve the DPO, Computer Centre, Legal Department, ethics committee, management of the relevant unit, or another specialist unit.

5. Cooperate on remediation. It may be necessary to invalidate a token, change a password, ask the provider to delete data, notify the system administrator, restrict access, or assess reporting obligations.

5. Ethics, rights of persons, and synthetic media

5.1 Ethics is not an add-on, but a basic condition

AI can affect trust, reputation, academic integrity, relationships between teachers and students, fairness of assessment, and the safety of individuals. Its use must therefore be proportionate, transparent, fair, and respectful of people's rights.

Before using AI, ask yourself three questions:

1. Is the use of the tool fair to the persons concerned?

2. Is it proportionate to the purpose and risk?

3. Can I explain the use of AI to a student, colleague, research participant, supervisor, supervisory authority, or the public?

5.2 Deepfakes, deep nudes, voice clones, and forged content

At UP, it is unacceptable to use AI to create or disseminate content that, without a legitimate reason and consent, imitates a specific person in a way that may damage their dignity, privacy, reputation, safety, or rights.

The following is prohibited or unacceptable in particular:

- creating sexually explicit synthetic materials depicting real persons,
- creating false voice recordings, videos, or photographs of real persons without a legitimate purpose and consent,
- creating forged evidence, records, signatures, images, or audio materials,
- disseminating synthetic media that may create the false impression that a person actually said or did something,
- using synthetic media for bullying, extortion, disinformation, reputational harm, or manipulation.

Exceptions may include clearly labelled educational, research, security, or awareness-raising demonstrations, provided they are proportionate, legally and ethically assessed, use synthetic materials or materials covered by consent, and minimise the risk of misuse.

5.3 Images, audio, video, and graphic outputs

AI rules do not apply only to text. For images, graphs, presentations, voice, music, video, 3D outputs, code, and datasets, always assess in particular authenticity, rights of persons, licence restrictions, suitability for publication, and the need to label the content as AI-generated or AI-assisted.

5.4 Copyright and licence restrictions

Do not enter third-party copyrighted texts, photographs, databases, recordings, textbooks, study materials, articles, or other works into AI tools to an extent that would conflict with the licence, contract, publisher's terms, or statutory exceptions. This also applies if the material is easily available online.

An AI output is not necessarily legally unproblematic simply because it was generated by an AI tool. Before publication or use in teaching, research, or promotion, it is necessary to assess in particular the origin of inputs, similarity to existing works, rights of persons, the tool licence, and the purpose of use.

6. Risky use of AI at the university

6.1 Prohibited or high-risk areas

Without prior legal, ethical, security, and methodological assessment, do not use AI to:

- determine access to studies or assign persons to educational programmes,
- perform automated assessment of study results if the output affects the course of education or the student's rights,
- monitor and detect prohibited student behaviour during tests,

- make decisions on recruitment, selection, performance assessment, promotion, or termination of employment,
- monitor the performance of employees or students,
- recognise emotions in an educational or workplace environment,
- carry out biometric identification, categorisation, or profiling,
- make decisions about the rights, obligations, or significant interests of individuals.

Use of AI in the above areas requires assessment of purpose, legal basis, proportionality, human oversight, documentation, security, transparency, and the possibility of appeal or remedy.

6.2 Human oversight and the option not to use the output

In higher-risk situations, it must be clear who is responsible for the decision, how the AI output is verified, and how the output can be rejected, corrected, or reversed. An AI output must not be used as an invisible authoritative statement, especially if it concerns a student, applicant, employee, research participant, or another identifiable person.

7. AI in teaching

7.1 Basic rule for teachers

The teacher has the right and duty to determine how AI is permitted, restricted, or prohibited in a given course, assignment, or examination. The rules must be communicated to students clearly and in good time. It is best to state the rules directly in the instructions for the specific activity rather than burdening the syllabus with extensive universal text.

A brief framework provision in the syllabus is sufficient, for example:

“The use of AI is governed by the rules stated for individual assignments and examinations. If no rules are stated, the student may use AI only as a support tool for language, technical, or organisational assistance; the student must not present AI-generated content as their own work and must disclose deliberate use of AI.”

More detailed rules should be included in the assignment instructions for a seminar paper, project, test, laboratory report, presentation, or examination.

7.2 Five-level model of AI use in a specific activity

For clarity, a five-level model may be used for individual activities. It is recommended to use it at the level of a specific assignment, not mechanically for an entire course.

Level	Mode	Appropriate use
0 – No AI	AI is not permitted.	A test of memorised knowledge, independent argumentation without aids, an assessment designed to verify individual competence.
1 – Technical and language assistance	AI may help with spelling, style, formatting, or translation, but must not change the substantive core.	Language editing of one’s own text, checking comprehensibility, typos, and format.
2 – Brainstorming and structure	AI may help with ideas, an outline, questions, or variants; the final content is created by the student.	Drafting an outline, generating counterarguments, preparing study questions.
3 – Assisted creation with disclosure	AI may create working parts of a text, code, visualisation, or analysis; the student must substantially revise and verify the AI-generated material and disclose the use of AI.	Programming, data analysis, graph drafts, first draft of an abstract.
4 – AI as part of the assignment	The use of AI is itself assessed; the student documents the workflow, prompts, verification, and reflection on limitations.	Critical evaluation of AI, model comparison, hallucination audit, prompt work, AI workflow design.

7.3 What rules for a specific assignment should contain

For every assignment in which AI may play a role, it is advisable to state:

- whether AI is permitted, restricted, or prohibited,
- for which activities the student may use it,
- whether the student must disclose the use of AI,
- how the student should describe the use of AI,
- whether prompts, outputs, or reflection must be attached,
- whether the process, the result, or both are assessed,
- whether specific types of assistance are prohibited, such as generating the entire text, solving tasks, rewriting third-party sources, or using non-public data.

7.4 Recommended wording for assignment instructions

Option A – AI prohibited for the substantive part

“Generative AI may not be used for the substantive solution of this assignment. Only basic checking of typos and formatting is permitted. The submitted text must be the result of the student’s own work.”

Option B – AI permitted as a support tool

“You may use AI for brainstorming, outlining, language editing, and questions to check understanding. You must not submit an unedited AI output as your own work. At the end, state which tool you used and for what purpose.”

Option C – AI as part of the work

“The use of AI is part of the assignment. Attach a brief description of the workflow, selected prompts, main outputs, the method of verification, and a reflection on the tool’s errors or limitations.”

7.5 AI detectors are weak evidence

Tools for detecting AI-generated text have limited reliability, especially for shorter texts, edited texts, texts in Czech, translations, and texts by non-native speakers. A detector may be only an indicative signal, not the sole evidence of a rule breach.

Where undisclosed use of AI is suspected, it is appropriate to combine several types of evidence:

- the assignment and its rules,
- work history, drafts, notes, and sources,
- oral defence or follow-up questions,
- comparison with the student’s usual performance,
- assessment of factual errors, citations, and methodology,
- communication with the student and faculty procedural rules.

7.6 Teach students to work responsibly with AI

The purpose of rules is not only to prohibit AI, but to teach students to recognise when AI helps and when it harms. Suitable teaching activities include:

- comparing an AI answer with an expert source,
- finding hallucinated citations,
- correcting a biased or one-sided output,
- reflecting on bias and errors in Czech,
- creating better prompts,
- documenting a workflow,
- discussing copyright, ethical, and data-related risks.

8. AI in science and research

8.1 AI as a research workflow tool

AI may be useful for preliminary literature searches, summarisation, work with literature, code creation, data analysis, annotation, visualisation, questionnaire preparation, language editing, translation, abstract drafting, or hypothesis generation. However, it is necessary to distinguish between preliminary assistance and a verified scientific output.

AI must not replace:

- the author's responsibility for data and interpretation,
- methodological justification,
- critical work with sources,
- peer review,
- research participants' consent,
- ethics-committee rules,
- grant and publication conditions.

8.2 Literature searches and citations

An LLM is not a bibliographic database. It may help with orientation in a topic, suggesting keywords, or explaining terms, but it cannot be relied on to provide existing sources or accurate citations.

When working with literature:

- verify the existence of every citation in a trustworthy database or library system,
- do not copy AI-generated references without checking them,
- distinguish between a summary of a source and your own assessment,
- do not enter non-public manuscripts or review materials into public tools,
- respect the licence terms of databases, publishers, and library resources.

8.3 Research data

Unpublished research data, data from clinical, psychological, educational, sociological, biological, or other research, and data obtained from respondents or research participants generally belong to the yellow or red zone. Before using AI, it is necessary to take into account informed consent, ethics approval, the data management plan, contractual obligations, grant rules, and the risk of re-identification.

Anonymisation must be genuine and proportionate. Merely replacing a name with a code often means only pseudonymisation, not anonymisation.

8.4 Transparency in scientific outputs

If AI has significantly assisted with text, analysis, code, visualisation, translation, or another part of the research process, it is advisable to describe its use in the methodology section, acknowledgements, or a note, according to the rules of the discipline, publisher, or grant agency.

Example wording:

"During preparation of the language version of the text, the tool [tool/model name] was used for stylistic checking and suggestions for wording. All substantive conclusions, citations, and interpretations were verified by the authors."

Example for data analysis:

"The tool [name] was used to draft part of the analytical script. The code was checked, modified, and validated by the authors on test data."

8.5 Review and confidential materials

A reviewer, editor, committee member, or evaluator must not enter a non-public manuscript, grant application, habilitation material, personnel file, non-public review, or other confidential material into a public AI tool. Use of AI

in these processes is possible only if the rules of the process expressly permit it and a secure processing mode exists.

9. AI in administration and support activities

9.1 Common administrative use

AI may help draft the structure of an e-mail, convert notes into minutes, prepare a document outline, edit the language of a public text, draft a presentation, summarise public materials, or prepare alternative wordings. In these cases, factual accuracy must be verified and the output adapted to the UP context.

9.2 Internal documents and contracts

Do not enter internal documents, contracts, economic data, payroll data, HR records, strategic documents, or non-public minutes into public AI tools. If AI is to assist with their processing, the use must take place in a permitted mode and the nature of the data must be taken into account.

Without consultation, do not use AI for:

- legal interpretation of a contract,
- decisions on legal actions in employment matters,
- processing HR materials,
- automatic evaluation of applicants,
- processing complaints, disciplinary or other sensitive submissions,
- generating responses in situations with legal impact.

9.3 Communication on behalf of UP

AI outputs intended for public communication, media, social networks, the web, press releases, marketing texts, or representation of UP must undergo human review. Factual accuracy, tone, alignment with the communication strategy, rights to the visual materials used, and the risk of reputational harm must be checked.

10. Practical scenarios

10.1 Public text

Situation: I want to improve the style of a public text about a conference.

Procedure: If the text contains no internal information, personal data, non-public data, or protected content, an ordinary tool may be used. Check and revise the output.

10.2 Meeting minutes

Situation: I want to enter minutes from an internal meeting into AI and ask for a summary.

Procedure: First remove personal, strategic, contractual, and sensitive information. If the minutes contain internal information, use only a permitted mode for the yellow zone. If sensitive information is involved, seek advice on how to proceed.

10.3 Student list

Situation: I want AI to create an analysis based on a list of students and test results.

Procedure: Do not use a public AI tool. These are personal data and study results. Use an approved internal procedure or consult the DPO and the responsible unit.

10.4 Non-public examination assignment

Situation: I want to use AI to generate variants of test questions from a non-public assignment.

Procedure: Non-public examination assignments do not belong in public tools. It is possible to work with a general description of the curriculum or with public sample questions, or with an internal permitted mode.

10.5 Grant application or manuscript

Situation: I want to enter a draft grant application or article into AI for language checking before submission.

Procedure: Check whether it contains non-public research know-how, personal data, sensitive data, or contractual restrictions. Do not use a public tool without contractual protection for confidential materials.

10.6 API key in a prompt

Situation: I accidentally entered an API key into an AI chat.

Procedure: Immediately stop working with the tool, invalidate the key or ask the administrator to invalidate it, change related secret values, record the circumstances, and contact the Computer Centre.

10.7 Suspected AI use in a seminar paper

Situation: A detector marked a seminar paper as probably created by AI.

Procedure: Do not use the detector as the only evidence. Check the assignment and rules, ask the student to explain the process, verify sources and factual errors, and, where appropriate, use an oral defence or a process under faculty rules.

11. Recommended process for new AI solutions at UP

Before introducing a new AI solution for teaching, research, administration, or support, it is advisable to carry out a short assessment:

1. Security: How are access, logs, authentication, encryption, and audit handled?

2. Data: What data will be entered and generated?

3. Control: How will the quality of outputs be verified, and how can an output be corrected or rejected?

4. Responsibility: Who is the process owner and who are the specialist units or officers sharing responsibility (for example, DPO, Computer Centre, Legal Department, ethics committee)?

5. Risk: Does the tool concern persons, assessment, decision-making, employment-law matters, research data, or public communication?

6. Contracts: What are the conditions for data processing, retention, and training?

7. Training: What minimum instruction do users need?

8. Transparency: How will students, employees, research participants, or the public be informed?

9. Purpose: What is the tool intended to be used for?

10. Users: Who will work with it?

12. Triage, specialist consultations, and responsibilities

If you are unsure or dealing with an incident, do not try to solve the problem merely by asking AI another question. Contact your supervisor or the responsible person at your workplace and the Vice-Rector for Research, Creative Activities and Knowledge Transfer. The Vice-Rector will perform triage and involve a relevant specialist unit according to the nature of the matter.

Topic	Procedure
Personal data, GDPR, DPIA, personal-data incident	Vice-Rector for Research, Creative Activities and Knowledge Transfer; DPO.
Contracts, licences, legal interpretation, liability	Vice-Rector for Research, Creative Activities and Knowledge Transfer; Legal Department.
Cybersecurity, accounts, API keys, access, technical configuration	Computer Centre or the relevant system administrator; in the event of a security incident, proceed according to information-security rules.
Academic integrity, assessment rules, student administration	Course guarantor, head of unit, study office, or faculty ethics committee.
Research ethics, research participants, sensitive research projects	Ethics committee or responsible unit.

Public communication, media, visuals, representation of UP	Communication Department or designated unit.

13. Practical glossary

AI agent – a system that, based on a goal, can plan steps, use tools, and perform actions, such as searching, reading files, writing e-mails, or running scripts.

AI literacy – the ability to understand the basic properties, limitations, risks, and appropriate ways of using AI.

AI hallucination – a situation in which AI creates an untrue, inaccurate, or fabricated output, often worded very convincingly.

Anonymisation – irreversible removal of the link to a specific person; data modified in this way are no longer personal data. Merely replacing a name with a code is usually not sufficient.

API – an interface through which software can communicate with another service, for example with an AI model.

API key / access token – a non-public access credential enabling the use of a service or account. It belongs to the red zone, similar to a password.

Autonomous decision-making – decision-making in which a system produces an output with an impact on a person without appropriate human assessment.

Bias – systematic distortion in inputs, data, the model, interpretation, or use of the output.

Deepfake – synthetic or modified image, audio, or video content that creates the impression that a specific person said or did something they did not in fact say or do.

Deep nudes – synthetic sexually explicit depiction of a person, typically created without their consent. Such use is unacceptable in a university environment.

DPA / data-processing agreement – a contractual arrangement between a controller and a processor of personal data.

DPIA / data protection impact assessment – an assessment of risks associated with planned processing where it may affect the rights and freedoms of persons.

Embedding – a numerical representation of text, an image, or other content, used for example for similarity search or RAG.

Generative AI – an AI system that creates new content, such as text, images, audio, video, code, tables, or proposed solutions.

High-risk AI system – an AI system falling into higher-risk areas, such as education, employment, biometrics, or significant decision-making about persons, where the legal framework requires special measures.

Human in the loop – a setting in which a human reviews the AI output and bears responsibility for its use.

LLM – a large language model, i.e. a type of AI model intended primarily for working with language.

Local model – an AI model operated on one's own device or within internal infrastructure; it may reduce some data risks but does not, by itself, ensure compliance with legal, security, and quality requirements.

No-training mode – a setting or contractual condition under which the provider does not use the user's inputs and outputs for further model training.

Opt-out – a setting that limits the use of data for service improvement. By itself, it does not replace contractual and security guarantees.

Prompt – an instruction, request, or context entered into an AI tool.

Prompt injection – an attack or unintended phenomenon in which text in a document, e-mail, web page, or other input instructs AI to behave in a way the user did not intend.

Pseudonymisation – replacement of direct identifiers with a code or placeholder. The data may still be personal data.

RAG / retrieval-augmented generation – a procedure in which AI answers using retrieved documents or a database. Quality depends on the quality of sources, access permissions for those sources, and correct configuration.

Data retention – the period and method of storing inputs, outputs, and logs by the provider or operator.

Approved / permitted mode of use – a specific combination of tool, purpose, data, contractual terms, security configuration, and responsibilities that the university or relevant unit has permitted.

Personal data controller – the entity that determines the purpose and means of personal data processing; in typical work situations at UP, this is the university.

Synthetic media – image, audio, video, or other media content created or substantially modified by AI.

Token as a model unit – a technical unit of processed text, approximately part of a word or a word; it is used for input and output length limits or service billing.

Incident triage – the initial assessment of reported or detected unauthorised AI use to determine its severity, urgency, and the appropriate procedure for handling it (whether it is a minor error, an academic offence, a security problem, a data-protection breach, or a serious ethical/procedural incident).

Public AI tool – a service available without special contractual protection or university approval; it is not suitable for non-public and sensitive data.

Zero Data Retention – a mode in which the provider does not retain content after processing, or retains it only minimally and in a technically defined manner.

Personal data processor – an external entity that processes personal data for the controller and according to the controller's instructions.

During work on this document, GPT-5.5 Thinking was used in full-collaboration mode to incorporate comments from active participants in the survey. GTP-6 Astra was used for the translation into English. The Vice-Rector for Research, Creative Activities and Knowledge Transfer, Drábek, thanks colleagues for their cooperation (in alphabetical order, without titles): Cahová, Černík, Dvořáček, Fürst, Hamuľáková, Hulicius, Chamrád, Kopecký, Kryštof, Laštovička, Látal, Mašát, Mazalová, Nétek, Olšan, Otřísal, Petr, Petřík, Šindlerová, Šubová, Tomášek, Tomášková, Vodák, Voráč, Walek, Změlík.