



UNIVERSITÀ
DEGLI STUDI
DI PADOVA

**Instructions for the use
of generative AI tools
in teaching and in
writing dissertations
and doctoral theses**

Contents

Introduction	3
AI at glance	3
What are language models and how do they work?	3
Introduction to chatbots	4
Chatbot limitations	5
AI in university teaching	6
Reducing the internal skills gap	6
Fostering reflection on one's own work	6
Improving the quality of work	6
Greater speed in carrying out tasks	7
Personalising education and tutoring	7
Fostering interdisciplinary dialogue	7
Increasing productivity	7
Tool proposed for the use of AI in individual course units	8
Learning with AI	11
AI applications in the study and preparation of materials	11
Designing and producing new materials	11
Studying	12
Text processing and revision	13
Exploring and searching for information	13
Data analysis and visualisation	13
Transparency on the use of AI in assignments and dissertations	14
Acknowledgements	15

Edited by the University Teaching Committee

Working group:

Fabio Grigenti - Department of Philosophy, Sociology, Pedagogy, and Applied Psychology

Carlo Mariconda - Department of Mathematics "Tullio Levi-Civita"

Sara Mondini - Department of Philosophy, Sociology, Pedagogy and Applied Psychology

Pier Paolo Paulesu - Department of Public, International and EU Law

Giorgio Satta - Department of Information Engineering

Valid from September 2025

Introduction

This document supplements and expands upon the “**Guidelines for the use of generative AI tools in teaching and in writing dissertations and doctoral theses**” offering practical advice, examples, and user experiences to support their application.

The instructions in this document are subject to updates, given the rapid evolution of these technologies/ methodologies and depending on any agreements between Unipd and generative AI providers.

Note: we have used ChatGPT4o (March 2025) to check for any unnecessary repetitions.

AI at glance

What are language models and how do they work?

At the heart of a chatbot, but also of many other generative AI systems, is a mathematical model called a “large language model” or LLM for short. Once the initial part of a sentence, which we will call the prefix, is given to an LLM, it calculates for each word in the vocabulary the probability that it will appear as a continuation of the prefix itself. For example, given the prefix “the book is on”, a good LLM should calculate relatively high probabilities for words such as “table” and “desk”, and low probabilities for words such as “jump” and “eat”. This is consistent with the fact that, according to our knowledge of the English language and the real world, words such as “table” and “desk” are a plausible continuation for the prefix “the book is on”.

On the other hand, the words “jump” and “eat” do not seem like a natural continuation for that prefix, as the resulting sentence would not have a clear meaning, or would even violate the grammar rules of English. There are two important stages in the life cycle of an

LLM. The first is called the training stage. In this stage, the LLM is trained on the basis of a vast amount of text retrieved from the web.

In this way, an LLM learns to statistically estimate the probabilities of words in a given context (prefix), processing syntactic and semantic

relationships. It must be noted that for an LLM, the context is not just a single sentence, but it covers several thousand preceding words. This means that an LLM can also capture the meaning of an entire discourse. The second stage in the life cycle of an LLM is the deployment stage. Here, the statistical information calculated by the model during the learning stage is “frozen”, and the model is used to generate sentences from a given prefix. This implies that, as a rule, the information passed onto an LLM in the form of a prefix has no influence on the model’s prior knowledge.

Over the last decade, the application of deep neural networks has made it possible to create LLMs that are capable of producing fluent and semantically coherent texts. This achievement has given a fundamental boost to the emergence of a new field of AI called generative AI.

Introduction to chatbots

An LLM is used to design a chatbot. As previously noted, an LLM is a tool capable of generating plausible texts from a prefix, but it is not explicitly designed to answer user questions in a comprehensive manner. For example, if you ask an LLM the question “What is the capital of France?” in prefix form, the model could most likely continue the sentence with the words “What is the capital of Germany?”, simply because the second question is very plausible in the context of the first.

It is therefore necessary to “force” the LLM to work in question-answer mode. To do this, it is sufficient to extend the training stage of an LLM by exposing the model to texts in the question-answer format covering a wide range of topics.

This way, our LLM essentially acts as a chatbot, i.e., it always interprets the input prefix as a question and generates a response deemed appropriate based on prior knowledge.

Given the length of the prefixes processed by an LLM, the questions that the user can ask a chatbot can also be very detailed and/or complex, such as requesting a summary of an entire document or generating original and/or creative text based on the general knowledge acquired during training. In order to control the quality of the responses, a technique commonly used in the LLM training stage is reinforcement learning which is described below.

Given a question to answer, a supervisor (typically human) expresses an opinion on the response received.

In the case of responses deemed inappropriate, the training algorithm

modifies the parameters of the neural network that constitutes the LLM in order to reduce the probability of inappropriate responses.

To fully exploit the potential of chatbots, it is important to communicate with them effectively through “prompts”.

Here are some tips: write clear, direct and complete prompts; avoid vague and ambiguous questions; specify the type of communication (e.g. formal or informal) and style; specify the length and format of the response you want to obtain; suggest a specific “role” for the chatbot to take on when generating the response.

For example, a text can be summarised by a chatbot in very different ways depending on whether the desired length, terminology and level of detail of certain concepts are specified. A response can be substantially different if we suggest the chatbot to take on the role of a professional technician or that of a populariser. Finally, replying to the chatbot to obtain further clarification or to refine the response also helps.

The University has made the Lucrez-IA chatbot, based on Anthropic’s Claude, available both on the web for university teachers with SSO access and on the various Moodle platforms to provide students with information about individual course units following instructions from the university teacher.

Chatbot limitations

It is important to note that the probabilities calculated by an LLM depend almost entirely on the texts that are read during the learning stage. As a result, any bias inherent in these texts is captured and inherited by the model. Research in the field of generative AI is experimenting with methods to minimise these effects but, given the enormous amount of data required during the training stage to obtain good probability estimates, results are not yet satisfactory.

Although the responses produced by a chatbot are fluid and plausible, we must bear in mind that they are not always the result of knowledge stored by the algorithm, but result from statistical properties derived from the texts read, as already described in relation to LLMs.

It is therefore possible that, however plausible a response may appear to the user, it is simply the result of the model’s invention! To this end, we must remember that an LLM is essentially a model trained to generate, even creatively, texts that are statistically very similar to those analysed during the training stage. Despite efforts to refine LLM training algorithms, there is currently no guarantee that the texts produced are

factually correct. It is therefore necessary to always carefully examine the texts produced by an LLM before using them.

AI in university teaching

Reducing the internal skills gap

Some studies show that lower-skilled workers benefit more from the use of AI tools than higher-skilled workers ([Peng et al., 2023](#); [Kanazawa et al., 2022](#)). This could reduce the skills gap between students with different levels of experience and competences. For example, a student with less writing experience can use AI tools to improve how their original results are presented.

Fostering reflection on one's own work

As a dialogical but artificial partner, AI tools can offer an external and detached perspective on one's own beliefs, arguments and ideas.

This is particularly useful in teaching and in academic writing, where an alternative point of view can stimulate critical thinking and selfassessment.

Improving the quality of work

Observations from other sectors indicate that the use of AI tools can aid in performing certain tasks to a higher standard.

In education, we expect university teachers and students to be able to achieve their goals with greater accuracy and depth. For example, a literature review can become more comprehensive and be carried out faster. When writing dissertations, AI can support linguistic and structural revision, helping to improve the clarity and coherence of the text and encouraging writing in languages other than one's mother tongue.

Again, AI can provide support to already solid linguistic knowledge, pointing out minor errors and typos. In any case, translation into languages that are not already mastered to a sufficient level is not recommended.

Greater speed in carrying out tasks

In today's academic environment, university teachers and students often face challenges that require new skills, such as writing projects, drafting management plans, and complying with data protection regulations and ethical guidelines.

The use of AI tools can certainly speed up the completion of tasks assigned.

Personalising education and tutoring

AI tools will soon make it possible to provide personalised responses to students via trained chatbots which will be able to take into account gaps in knowledge and progress made.

As of February 2025, this is already partially possible on Moodle course unit pages via Lucrez-IA, with a limit on the number of characters in the instructions given by the university teacher.

Fostering interdisciplinary dialogue

AI tools facilitate the exploration and understanding of a vast amount of information, making it easier to link academic work from different disciplines. This can foster interdisciplinary collaboration by opening up new learning opportunities ([Chubb, Cowling & Reed, 2022](#)).

Increasing productivity

Productivity can be defined as the amount of desired results obtained using a given amount of resources. Scientific studies show that the use of AI tools can significantly increase productivity ([Brynjolfsson et al., 2023](#); [Noy & Zhang, 2023](#); [Al Naqbi et al., 2024](#)). In the near future, tasks that require time-consuming effort, which do not require significant human input, may be partially delegated to AI tools. These resources can also enhance users' skills (e.g., by assisting them in creative development). In terms of teaching and writing dissertations, AI can speed up collecting information and organising content, allowing more time for critical reflection, writing the final draft, and preparing work material.

Overall, we expect these benefits for individual university teachers and students to translate into a greater social impact of education with the same resources available. This implies that the resources used in education will be used more effectively, potentially leading to more significant results.

Tool proposed for the use of AI in individual course units

We propose a tool (Table 1) to provide teaching staff with an optional list of instructions on the use of AI in individual course units.

These instructions are aimed at both students and staff involved in teaching. The tool identifies a number of independent aspects relating to the use of AI in teaching.

For each aspect, three different approaches are described:

- allowed;
- conditional or restrictive;
- not allowed.

University teachers can therefore choose the best combination of aspects for their course unit and inform students and collaborators of the instructions.

This tool can be used without modification or adapted to the specific needs of the course unit by modifying the text and/or adding new items as appropriate. In addition, university teachers are invited to open up opportunities for debate on the topic tailored to the specific needs of the Academic Discipline (SSD), Degree Course, Department and School, without in any way hindering the autonomy of teachers, but acting with a strategic vision.

NOTE ON METHODOLOGY USED

The configuration tool described here was created by gathering best practices from international universities ([Vanderbilt](#), [MIT](#), [Harvard](#), [NY University](#), [Georgia Tech](#), [Texas-Austin](#), [Penn State](#)), classifying the various recommendations and translating them into Italian.

Subsequently, we interacted in English with ChatGPT4o (plus version, September 2024), starting with the prompt: “Artificial intelligence may be used in higher education. What alternatives can be reported in a syllabus (e.g., admitted during exams/used in the course, etc.)?”, and including the documents collected in the prompt. In particular, areas of AI use were identified for learning and skills development; academic collaboration and engagement; feedback and assessment.

Finally, the text produced was translated into Italian and approximately 20% of the content was manually reworked.

SCOPE OF USE	POSITION
As a teaching tool	<input type="checkbox"/> <i>Allowed:</i> Generative AI can be used by teaching staff to improve the teaching experience (e.g. chatbots, tutors, virtual assistants, simulations, data analysis). <input type="checkbox"/> <i>Conditional:</i> The use of generative AI as a teaching tool may be limited to specific activities or parts of the course. <input type="checkbox"/> <i>Not envisaged:</i> Generative AI will not be used as a teaching tool in the course, as it may interfere with the development of specific skills, as set out by the university teacher's criteria
For the development of technical skills (e.g. GitHub Copilot for programming, chatBot for data analysis)	<input type="checkbox"/> <i>Allowed:</i> Students may use assistants and other generative AI tools to develop technical skills. <input type="checkbox"/> <i>Conditional:</i> Students may use virtual assistants and other generative AI tools to develop technical skills, but must demonstrate they understand on their own the results generated by such tools. Students must cite the tools used and include the relevant prompts. <input type="checkbox"/> <i>Not allowed:</i> Based on the learning outcomes, the university teacher does not allow the use of generative AI to develop technical skills
For assignments and essays	<input type="checkbox"/> <i>Allowed:</i> Students may use generative AI tools to facilitate peer learning brainstorming, group work, or peer reviews. <input type="checkbox"/> <i>Conditional:</i> Generative AI tools may only be used to support discussions, but not to replace student contributions. Students must cite the generative AI tools used and summarise how and where they were used. <input type="checkbox"/> <i>Not allowed:</i> Generative AI may not be used to replace collaborative activity, discussion or communication within a group.
For collaborative and peer learning	<input type="checkbox"/> <i>Allowed:</i> Students may use generative AI tools to facilitate brainstorming, group work, or peer reviews. <input type="checkbox"/> <i>Conditional:</i> Generative AI tools may only be used to support discussions, but not to replace student contributions. Students must cite the generative AI tools used and summarise how and where they were used. <input type="checkbox"/> <i>Not allowed:</i> Generative AI may not be used to replace collaborative activity, discussion or communication within a group.
For accessibility	<input type="checkbox"/> <i>Allowed:</i> Both the university teacher and students may adopt generative AI tools that improve accessibility (e.g. text-to-speech, machine translation) to foster inclusion. <input type="checkbox"/> <i>Conditional:</i> The use of generative AI tools is only allowed for students with specific needs (e.g. disabilities, language barriers).
For feedback and assessment	<input type="checkbox"/> <i>Conditional:</i> The use of generative AI is limited to feedback on low-stakes assignments and educational assessment in general, while final assessment and overall assessment in general will be handled by the university teacher. <input type="checkbox"/> <i>Not allowed:</i> The use of generative AI tools for student assessment is not allowed.

SCOPE OF USE	POSITION
During examinations	<input type="checkbox"/> <i>Conditional:</i> The use of certain generative AI tools may be permitted under certain conditions, e.g. open-book examinations or examinations for specific purposes. <input type="checkbox"/> <i>Not allowed:</i> To ensure academic integrity, generative AI tools are not permitted during examinations.
For academic writing	<input type="checkbox"/> <i>Allowed:</i> Students may use generative AI tools to manage citations, look for sources, or organise research, but their writing must be original. <input type="checkbox"/> <i>Conditional:</i> Students may use generative AI only for preliminary activities (e.g. creating bibliographies), but the final analysis must be their own. <input type="checkbox"/> <i>Not allowed:</i> The university teacher deems that, at the stage of learning in which the teaching takes place, the use of generative AI tools for academic writing activities is not suitable and therefore not allowed.

Learning with AI

The main AI-based applications in the fields of teaching and dissertation and theses writing can be grouped into five macro areas. Each area has its own specific issues. For general issues, please refer to the General Principles of the “**Guidelines for the use of generative AI tools in teaching and in writing dissertations and doctoral theses**”.

AI applications in the study and preparation of materials

Designing and producing new materials

This area includes the use of AI in the production of original material.

More specifically:

1. criticism of hypotheses;
2. brainstorming;
3. producing ideas;
4. generating texts based on a user-suggested outline;
5. generating code;
6. generating summary data.

Criticism made by AI should not necessarily be taken as irrefutable argument, but rather as a starting point for further debate (point 1). Ideas formulated and perceived as new by AI should always be reviewed to verify their accuracy and sources, so as to avoid the risk of presenting other people’s ideas as one’s own without citing them. AI may not have taken recent developments into account or may not have had access to all the most recent literature prior to the time it is used by the user (points 3–4). The same applies to any code generated, which must always be checked and tested before being included in one’s scientific output (point 5). The generation of summary data must always be reported in detail, because, if omitted, it can lead to fraudulent research.

Studying

In addition to producing documents, AI can be used by students as an assistant for individual study to aid them in every phase of learning, with the usual recommendation to check the output of AI against other methods (see the MOOC “[Imparare con l’IA](#)” by S. Sancassani and D. Casiraghi - POK Polimi).

Exploration:

- Deepening and understanding concepts, also through analogies
- Analysing the structure of a lesson or text
- Collecting sources
- Using diagrams and images for complex concepts

Revision:

- Creating a glossary • Creating guiding questions

Application:

- Reviewing key concepts
- Breaking down a complex problem
- Highlighting frequent errors
- Creating a summary of errors and strategies

Understanding through classical dialectics:

- Discussing a topic through thesis/definition/argumentation/ deductio n/rebuttal/synthesis
- Identifying and analysing logical fallacies
- Interpreting roles and perspectives

Consolidation:

- Defining content
- Creating feedback and assessment questions
- Changing context and reversing roles
- Practising orally

Reflecting on one’s studies:

- Defining the aims and content of the assessment
- Asking questions and evaluating the answers
- Reflecting on one’s comprehension
- Analysing one’s study method
- Readjusting study strategies

Text processing and revision

This category includes applications that analyse the prompt typed by the user (or the file uploaded) and return an analysis of it.

More specifically:

1. *Summary* of texts and/or articles.
2. *Feedback* on scientific hypotheses.
3. *Structuring* the dissertation (e.g. production of a table of contents).
4. Linguistic and/or grammatical *revision*.

Some tips specific to this context include: citing only works consulted in their entirety in the sources and not those indicated or summarised by the AI (point 1); checking that the suggestion provided by the AI is consistent with the work in general (point 2); double-checking grammatical and/or linguistic suggestions each time (point 4).

Exploring and searching for information

This area includes all those applications whose main purpose is to provide information on literature relating to a specific field and/or working hypothesis. More specifically:

1. *Looking* for bibliographic information.
2. *Analysing* the content of the information found.
3. *Asking for a preliminary opinion* on a specific working hypothesis, based on the literature available.

For the correct use of AI in this specific area of application, first of all we recommend *always checking manually* the literature suggested by the AI, thus avoiding using its suggestions without first checking them (point 1); personally analysing the texts in their entirety before performing a more targeted automatic search on individual occurrences (point 2); not taking the preliminary opinion returned by the AI as proof, but only as an initial answer that always requires subsequent systematic review (point 3).

Data analysis and visualisation

This area includes all applications that return or enable the visualisation and analysis of data through

AI and/or, more in general, machine learning software.

The general advice in this case is to check that the initial dataset complies with ethical and legal principles (e.g. check for bias) and to check that the prompt has been properly understood.

Transparency on the use of AI in assignments and dissertations

Always declare the use of AI tools in submitted work, specifying the tool, version and, if required, the prompts used, using APA or MLA standards. Here are some examples of citations using the two standards.

APA style citation (7th edition)

Author's prompt: 'What are the main impacts of artificial intelligence in university education?'

Response generated by ChatGPT (OpenAI, 2025):

'Artificial intelligence is transforming university education through the personalisation of learning, the automation of administrative tasks, and access to advanced educational resources.'

Reference in bibliography:

OpenAI. (2025, 11 March). Response generated by ChatGPT to an author's prompt [GPT-4 model]. ChatGPT. <https://chat.openai.com>

In-text citation: (ChatGPT, OpenAI, 2025)

MLA style citation (9th edition)

Author's prompt: 'What are the main impacts of artificial intelligence in university education?'

Response generated by ChatGPT (OpenAI, 2025):

'Artificial intelligence is transforming university education through the personalisation of learning, the automation of administrative tasks, and access to advanced educational resources.'

Reference in bibliography:

OpenAI. Response generated by ChatGPT to an author's prompt [ChatGPT, version GPT-4, 11 March 2025, <https://chat.openai.com>].

In-text citation: (OpenAI)

Acknowledgements

Several committees contributed to discussions with members of this committee on the issue of guidelines and instructions for the use of generative AI and have been an inspiration for this document.

We would like to thank, in particular, Dr D. Da Re, Director of the Digital Learning Office, and our colleagues.

Professors G. Amore, A. Gentili, G. Da San Martino,
M. C. Lavagnolo, F. Mazzini, A. Menini, J. E. Raffaghelli, A. Rodà,
C. Sandei, F. Vespignani.

It should be noted that the choices made in this document are not necessarily attributable to the colleagues mentioned, who may not agree with all of its contents.

unipd.it