

INSIGHT: Bridging the Student-Teacher Gap in Times of Large Language Models

Jarne Thys*, Sebe Vanbrabant, Davy Vanacken and Gustavo Rovelo Ruiz

UHasselt – Hasselt University, Digital Future Lab – Flanders Make, Wetenschapspark 2, 3590 Diepenbeek, Belgium

Abstract

The rise of AI, especially Large Language Models, presents challenges and opportunities to integrate such technology into the classroom. AI has the potential to revolutionize education by helping teaching staff with various tasks, such as personalizing their teaching methods, but it also raises concerns, for example, about the degradation of student-teacher interactions and user privacy. Based on interviews with teaching staff, this paper introduces INSIGHT, a proof of concept to combine various AI tools to assist teaching staff and students in the process of solving exercises. INSIGHT has a modular design that allows it to be integrated into various higher education courses. We analyze students' questions to an LLM by extracting keywords, which we use to dynamically build an FAQ from students' questions and provide new insights for the teaching staff to use for more personalized face-to-face support. Future work could build upon INSIGHT by using the collected data to provide adaptive learning and adjust content based on student progress and learning styles to offer a more interactive and inclusive learning experience.

Keywords

AI Teaching Assistant, Teaching Support, Student-Teacher Interaction

1. Introduction

Due to the rising popularity of Artificial Intelligence (AI), various educational tools incorporating AI have been developed in recent years. AI can provide personalized learning experiences by tailoring educational experiences to individual learners' needs and abilities [1]. For example, adapting content and pace can improve learning effectiveness and efficiency [2, 3]. AI enables personalized learning at any time, making education more efficient and accessible to larger groups [4]. Although most research focuses on K-12 settings, personalized learning is also promising for higher education, which is the main focus of this work, as AI can consider the students' unique motivations and circumstances [3]. Recent research also focuses on making Large Language Models (LLMs) fit for the classroom by reducing hallucinations, which is achieved by using course materials as input for the answer [5]. In mathematics education, LLMs show promise in developing students' problem-solving skills and critical thinking abilities by enabling them to engage more deeply with open-ended problems through iterative questioning, clarification of concepts, and exploration of multiple solution paths [6]. Nitze [7] developed a prototype to simulate oral exams in STEM education to reduce educator workload and give students individual feedback in the early stages of their academic journey.

Despite the numerous benefits AI offers to students and teaching staff, its integration into education raises several concerns. First, AI can degrade human interaction in education [8, 9], as it reduces face-to-face engagement and creates excessive dependence on the technology [4, 10], which could lead to worse educational experiences for students. Second, as students start asking more questions to

D-SAIL Workshop - Transformative Curriculum Design: Digitalisation, Sustainability, and AI Literacy for 21st Century Learning, July 22, 2025, Palermo, Italy

*Corresponding author.

✉ jarne.thys@uhasselt.be (J. Thys); sebe.vanbrabant@uhasselt.be (S. Vanbrabant); davy.vanacken@uhasselt.be (D. Vanacken); gustavo.roveloruiz@uhasselt.be (G. Rovelo Ruiz)

🌐 <https://www.uhasselt.be/en/who-is-who/jarne-thys> (J. Thys); <https://www.uhasselt.be/en/who-is-who/sebe-vanbrabant> (S. Vanbrabant); <https://www.uhasselt.be/en/who-is-who/davy-vanacken> (D. Vanacken); <https://www.uhasselt.be/en/who-is-who/detail/gustavo-roveloruiz> (G. Rovelo Ruiz)

🆔 0009-0005-9348-7405 (J. Thys); 0009-0001-7996-6048 (S. Vanbrabant); 0000-0001-8436-5119 (D. Vanacken); 0000-0001-7580-8950 (G. Rovelo Ruiz)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

LLMs, the teaching staff might miss important cues, such as gaps in comprehension of the learning material [10, 9]. Third, privacy is an important concern, as AI tools collect a vast amount of data. OpenAI, for example, collects the user’s input content (including files), device information, and location information [11]. Educational applications, when deployed into a classroom, collect additional data that needs to be stored in a way that is both ethical [12] and compliant with applicable laws, such as the EU’s GDPR. Anthology’s BlackBoard, for example, collects, among other data, a student’s solutions, grades, progress on a per-file basis, average hours in the course, and days of inactivity [13].

To address these challenges, we introduce INSIGHT: INtelligent Student and Instructor Guidance through Human-Centered Technology, which implements different AI tools to support **student-teacher interaction** in a privacy-aware manner. Our contribution is twofold. First, INSIGHT equips teaching staff with **data-driven insights** on students’ questions and challenges regarding a specific course topic. Based on these insights, they can support each student in a tailored manner and adapt their course material based on the needs of students. Second, INSIGHT provides students access to an LLM in a monitored environment while ensuring **data privacy** through explicit opt-in data collection.

2. Interviews with teaching staff

To guide INSIGHT’s development, we organized focus groups with various teaching staff at Hasselt University. This was done in two phases. In the first phase, we conducted semi-structured interviews with members of the teaching staff from different courses. An initial prototype was built using this feedback. Afterward, we collected additional input by demonstrating the prototype at an internal workshop that specifically targeted AI in the university. This way, we could get input from a variety of people with different backgrounds.

The main takeaway from these interactions is that teaching staff are concerned they might lose contact with students, and thus, hinder their insights into the comprehension levels of their course knowledge. Traditionally, they could determine this based on how many and what kind of questions students would ask them. With the emergence of tools such as ChatGPT, they are concerned that students will start asking most of their questions to those kinds of tools, leaving the teaching staff with little feedback to improve their course. This was particularly relevant for a subset of the teaching staff who had only recently developed a new master’s program with completely new courses and course material.

3. INSIGHT’s modular framework for AI-powered student-teacher interaction

INSIGHT’s components are designed modularly, enabling INSIGHT to be deployed across various domains without changing its core implementation. We achieve this flexibility by dividing the solution into multiple independently adjustable components, as illustrated in Figure 1. We have a central system, *INSIGHT Core*, with a keyword extraction and sentence similarity model and a database. This database contains all the course information and data on student interactions with INSIGHT while solving exercises (e.g., questions to the LLM, FAQ usage, and difficulty ratings). INSIGHT Core is connected to a local LLM and a user interface that can be independently changed based on specific needs. Importantly, INSIGHT does not interfere with the behavior of the LLM in any way. There is no prompt engineering, rephrasing, or nudging of responses, ensuring students interact with it as they naturally would and are not discouraged by artificial limits. As a proof of concept, we incorporate data from Hasselt University’s Algorithms and data structures course [14] to demonstrate and assess INSIGHT’s feasibility. An overview of INSIGHT’s features and their benefits for teaching staff and students can be found in Table 1.

Answering student questions through LLMs. LLMs have become increasingly widespread, which presents a valuable opportunity for the educational sector to provide frameworks that guide students in

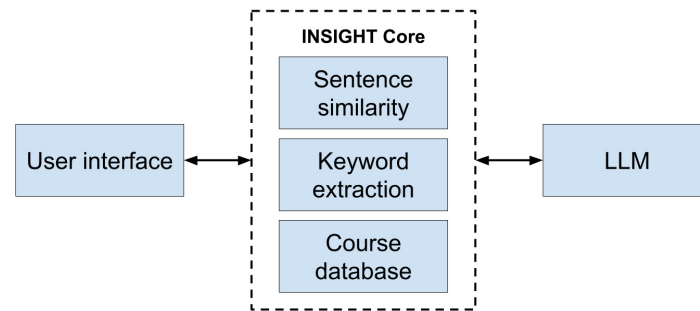


Figure 1: INSIGHT has a central system, *INSIGHT Core*, with a keyword extraction and sentence similarity model and a database containing all course information and student interactions with INSIGHT while solving exercises. INSIGHT Core is connected to a local LLM and a user interface that can be independently changed based on specific needs.

Feature	Benefit for teaching staff	Benefit for students
LLM chat interface	Reduction in workload as it can answer simpler questions	Easy access to answers, just like using other LLM tools
Dynamic FAQ	Allow verification of the answer to common questions	Have a teacher-verified answer for common questions
Teaching insights	Quick indication of topics of recurring questions	Get more personalized feedback, as the teacher knows where they might struggle
Data anonymization	/	Be in control of their data and have the option to share specific data they agree to

Table 1

Overview of the four INSIGHT features — an LLM-based chat interface, a dynamic FAQ generator, data-driven teaching insights, and a configurable data-anonymization layer — and the way each feature simultaneously benefits teaching staff (by reducing repetitive question load, enabling verification of common queries, and surfacing areas of recurring difficulty) and students (by providing instant, teacher-reviewed answers, more personalized feedback in the classroom, and the option to share only the data they consent to).

making the most of these technologies in a meaningful, ethical, and educationally responsible manner. We use an LLM as the main interaction method between students and INSIGHT. We later analyze students’ interactions with the LLM to provide the teaching staff with data they can use to offer better face-to-face support and increase student-teacher interaction. In this prototype, we used Llama 3.2 3B as it has an above-average output speed and below-average latency, ensuring quick responses [15]. This comes at the cost of intelligence, but in our initial testing, it seems adequate for the use case. The small amount of parameters also makes it easier to deploy on consumer hardware. Finally, INSIGHT’s modular architecture allows us to easily switch to other (newer) models and pick the right model depending on the needs of the course. Gemma 3 [16] is an example of a recent development that could be interesting for our use case. At the time of writing, it is a state-of-the-art model optimized to run on a single GPU. This would make it feasible to, for example, run one high-performance PC locally in the classroom.

Keyword mapping and dynamically generated FAQ from student questions. After the teaching staff review the keywords for each exercise, INSIGHT treats that list as the topic vocabulary for the course. All questions from students are embedded in the same vector space. The nearest keyword determines the question’s topic label. After identification, the topic is added to the analytics for the teaching staff. These student questions to the LLM can also help teaching staff identify challenging topics and knowledge gaps. Recurrent questions can be added to an FAQ, which is primarily designed as a tool for the teaching staff to monitor learning trends and gaps, but also serves as a resource for students facing similar issues. This new information also allows teaching staff to provide better face-to-face support for specific students or groups. To group similar questions in the FAQ, INSIGHT uses the vector

Select Keywords

Exercise 1	Exercise 2	Exercise 3
<input checked="" type="checkbox"/> boost libraries	<input checked="" type="checkbox"/> calculates median	<input checked="" type="checkbox"/> binary tree
<input checked="" type="checkbox"/> search value	<input type="checkbox"/> median ordered	<input type="checkbox"/> number tree
<input type="checkbox"/> boost	<input checked="" type="checkbox"/> array time	<input checked="" type="checkbox"/> tree node
<input checked="" type="checkbox"/> library java	<input checked="" type="checkbox"/> median	<input type="checkbox"/> tree tree
<input type="checkbox"/> available boost	<input checked="" type="checkbox"/> time complexity	<input type="checkbox"/> tree answer
<input checked="" type="checkbox"/> value array	<input checked="" type="checkbox"/> complexity terms	<input checked="" type="checkbox"/> tree
<input checked="" type="checkbox"/> array	<input checked="" type="checkbox"/> complexity	<input checked="" type="checkbox"/> binary
<input checked="" type="checkbox"/> search	<input checked="" type="checkbox"/> array	<input type="checkbox"/> zero
<input checked="" type="checkbox"/> library	<input checked="" type="checkbox"/> function	<input checked="" type="checkbox"/> number
<input checked="" type="checkbox"/> functions	<input checked="" type="checkbox"/> notation	<input checked="" type="checkbox"/> steps

Figure 2: The keyword selection is a mixed-initiative process where the initial list of keywords is determined through a keyword extraction algorithm. The teaching staff can further refine the list for final use. Table 2 shows the full texts for these three exercises, together with the automatically extracted keywords.

embeddings from the all-MiniLM-L6-v2 model to group semantically related questions based on their cosine similarity. This model was chosen because it was trained on multiple datasets, making it very generalizable, while it maintains a relatively high accuracy and speed [17]. We note that this approach was developed for the given algorithms and data structures and might require fine-tuning to apply it to other courses. Another benefit of grouping similar questions is the ability to return the same answer to these questions. INSIGHT caches the LLM’s answer and provides the teaching staff the ability to edit the answer to their liking, ensuring human-in-the-loop involvement and reducing the variability of LLM answers.

Leveraging student interaction data for teaching insights. When students engage with an LLM, teaching staff may lose sight of their understanding of the course material, as the amount and nature of questions serve as a form of implicit feedback on the learning resources. This makes it harder for the teaching staff to provide support and refine those learning resources. To address this, we propose a system that collects and analyzes students’ questions to the LLM, their FAQ usage, and explicit input in the form of difficulty ratings for the exercises they complete. This data provides valuable insights that, over time, can help refine course materials to better align with student needs and expectations. To determine the topics from students’ questions to the LLM, INSIGHT must first identify the topics covered in the exercises. To achieve this, we use keyword extraction to infer the exercise topics from the exercise descriptions. The keyword extraction is a mixed-initiative process that first uses KeyBERT [18] to extract an initial list of keywords from the exercise text. After extraction, the teaching staff can manually review the keywords for each exercise to best represent the tasks and make them suitable for identifying question topics. An example is illustrated in Figure 2.

Data anonymization policy. While it can be helpful for teaching staff to identify exercises and topics students find challenging, not all students might be comfortable sharing identifying data. To address this, students can select the *anonymous mode*, allowing teaching staff access to the data without the students’ names attached. Additionally, INSIGHT uses Ollama to run an LLM locally, ensuring student queries, potentially containing sensitive information, remain private and are never shared with third parties. This method ensures that personal data never leaves the students’ devices without their explicit consent.

User interface. The teaching staff’s user interface, illustrated in Figure 3, has two main components: a dynamic FAQ and two visualizations. The dynamic FAQ displays frequently asked student questions,

Exercise text	Extracted keywords
According to the documentation, which kind of functions do the C++ standard library and java provide to search a value in an array? What is available in the boost C++ libraries? Use these libraries (unless instructed otherwise) to make your work easier!	boost libraries, search value, boost, library java, available boost, value array, array, search, library, functions
This simple function calculates the median from an ordered array. Describe its time complexity in terms of the big o notation. What is the main difference here, compared to similar exercises on arrays when considering n?	calculates median, median ordered, array time, median, time complexity, complexity terms, complexity, array, function, notation
A binary tree is a tree in which each node has zero, one, or two children (see figure). How many steps would a search for a number take in such a tree? Answer in terms of n. Note that trees will be covered in later chapters.	binary tree, number tree, tree node, tree tree, tree answer, tree, binary, zero, number, steps

Table 2

Extracted keywords for three exercises from the Algorithms and data structures course. Through the mixed-initiative approach, the teaching staff can manually refine the keywords, as the AI model can make mistakes. An example can be seen in Exercise 3, where the model recommends 'tree tree' as a good keyword.

Algorithms and Data Structures

FAQ

Select an exercise: Exercise 1

Add FAQ

Q1: When should I use `std::find()` instead of `std::binary_search()` in C++?

Delete

Q2: What is the difference between `std::binary_search()` and `std::lower_bound()`?

Delete


Q3: Which is faster, `std::find()` or `std::unordered_set::find()` in C++?

Delete

Q4: Should I always use Boost libraries, or is the Standard Library sufficient?

Delete

FAQ Read Count



Subject Count




Figure 3: User interface for the teaching staff, with the dynamic FAQ and two visualizations. The teaching staff can edit the FAQ items and their content, or add an FAQ item manually. The visualizations show the frequency with which each FAQ item has been viewed and the frequency of topics in the questions from students to the LLM.

and staff can edit the LLM-generated answers to ensure accuracy. The two graphs visualize FAQ view frequencies and the frequency of course topics in student queries to the LLM. The student interface, illustrated in Figure 4, also includes the FAQ and an additional chat interface for the LLM.

4. Discussion

INSIGHT presents an approach to integrating AI in education to enhance student-teacher interaction while maintaining transparency in data collection and ensuring users retain control over their data. Its modular design allows for flexible integration in various courses, while its AI-driven FAQ and

Algorithms and Data Structures

Select an exercise:

FAQ

When should I use `std::find()` instead of `std::binary_search()` in C++?

What is the difference between `std::binary_search()` and `std::lower_bound()`?

Which is faster, `std::find()` or `std::unordered_set::find()` in C++?

Should I always use Boost libraries, or is the Standard Library sufficient?

Current Difficulty Rating: 3

Rate Exercise

Chat

You:
 How does linear search work?

INSIGHT:
 Linear Search is a basic algorithm used to find an element in a list or array by comparing it with each element one by one, starting from the first index.

 Here's how it works:

 1. **Initialization:** Start at the first index of the list.

Type your question here...

Send

Figure 4: The student interface includes the FAQ and a chat window for the LLM. FAQ views are counted, optionally attaching the student’s name with their explicit consent. Chat questions are analyzed to visualize question topics for teaching staff. Students can also rate exercise difficulty to provide explicit feedback.

keyword analysis help identify common student challenges. This knowledge can help teaching staff in personalizing face-to-face support for groups or individual students, depending on whether the individual student wants to share their data.

Despite these advantages, potential risks must be considered, including over-reliance on AI and the accuracy of generated responses. While the LLM’s capabilities and limitations with regard to answering students’ questions are not the focus of this work, advanced LLM prompting strategies, such as Chain-of-Thought [19], and the use of a larger LLM, could be worthwhile improvements.

Importantly, AI tools are now ubiquitous, and it is difficult, if not impossible, to control how students use them. Responsible AI use relies not only on system design but also on human practices like dialogue, trust-building, and setting shared expectations, elements that must be cultivated outside the boundaries of any technical system. Because of this, INSIGHT is not designed to impose hard constraints but instead tries to adopt an approach that encourages open conversations between students and teaching staff about how and when to use AI responsibly, using the FAQ system as a guide.

Furthermore, as students are aware that INSIGHT monitors their questions, they may ask questions to another LLM outside of INSIGHT. This can create a biased view for the teaching staff and skew their perception of students’ knowledge. To encourage the adoption of INSIGHT, students have to notice a positive impact on their learning experience using this tool, and thus, teaching staff need to actively use the new insights to improve student-teacher interactions. The teaching staff also has to communicate to the students when to use INSIGHT, as it is complementary to the staff and not a replacement. Certain questions are best answered by the teaching staff through direct, in-person interaction with students. Also note that INSIGHT fits within Hasselt University’s approach of interactive classes with teaching assistants, but may need to be revised to accommodate other class settings.

Finally, we want to emphasize that INSIGHT is not intended to prevent students from using LLMs to obtain direct answers. Instead, it creates an environment where students can use these tools with privacy safeguards in place, and where the resulting interaction data can be meaningfully leveraged by both students and teaching staff to enhance learning. While we recognize the value of scaffolding-based tools that guide students toward discovering answers themselves, INSIGHT addresses a different but equally important need: providing a safe and transparent space for students who primarily seek answers, ensuring that their use of LLMs still contributes to their learning process and informs face-to-face support.

The next step in our research is to empirically validate the usefulness of INSIGHT. To achieve this, we will first perform a cognitive walkthrough with new members of teaching staff who have expressed

interest in piloting the tool. Afterward, we will introduce INSIGHT as an assistive tool in various courses at Hasselt University. Future research opportunities lie in integrating knowledge tracing to support advanced personalization while providing more granular data on the uptake of specific skills to the teaching staff. Furthermore, future work could use the collected data to provide adaptive learning and adjust content dynamically based on student progress and learning styles, aiming to offer a more interactive and inclusive learning experience.

5. Conclusion

INSIGHT is our attempt to exploit the potential of AI-driven educational tools to enhance both student learning and teaching staff support. INSIGHT's goal is to help teaching staff identify knowledge gaps, personalize face-to-face support, and maintain student engagement by providing a modular, data-driven approach to analyzing student interactions. Its emphasis on privacy and transparency ensures that students can benefit from AI assistance without blindly giving away their data. In the next steps of our research, we will test INSIGHT's effectiveness in real-world courses to validate our approach.

Acknowledgments

This work was supported by the Special Research Fund (BOF) of Hasselt University (BOF24OWB28 and BOF23OWB31). This research was made possible with support from the MAXVR-INFRA project, a scalable and flexible infrastructure that facilitates the transition to digital-physical work environments. The MAXVR-INFRA project is funded by the European Union - NextGenerationEU and the Flemish Government.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly in order to: Improve writing style, Grammar and spelling check. After using these tool(s)/service(s), the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] L. Chen, P. Chen, Z. Lin, Artificial Intelligence in Education: A Review, *IEEE Access* 8 (2020) 75264–75278. URL: <https://ieeexplore.ieee.org/document/9069875/>. doi:10.1109/ACCESS.2020.2988510.
- [2] A. Shemshack, J. M. Spector, A systematic literature review of personalized learning terms, *Smart Learning Environments* 7 (2020) 33. URL: <https://slejournal.springeropen.com/articles/10.1186/s40561-020-00140-9>. doi:10.1186/s40561-020-00140-9.
- [3] K. Wozniak, Personalized Learning for Adults: An Emerging Andragogy, in: S. Yu, M. Ally, A. Tsinakos (Eds.), *Emerging Technologies and Pedagogies in the Curriculum*, Springer Singapore, Singapore, 2020, pp. 185–198. URL: http://link.springer.com/10.1007/978-981-15-0618-5_11. doi:10.1007/978-981-15-0618-5_11, series Title: Bridging Human and Machine: Future Education with Intelligence.
- [4] Z. Altinay, F. Altinay, G. Dagli, R. Shadiev, A. Othman, Factors Influencing AI Learning Motivation and Personalisation Among Pre-service Teachers in Higher Education, *MIER Journal of Educational Studies Trends and Practices* (2024) 462–481. URL: <https://www.mierjs.in/index.php/mjestp/article/view/2714>. doi:10.52634/mier/2024/v14/i2/2714.
- [5] K. Taneja, P. Maiti, S. Kakar, P. Guruprasad, S. Rao, A. K. Goel, Jill Watson: A Virtual Teaching Assistant Powered by ChatGPT, in: A. M. Olney, I.-A. Chounta, Z. Liu, O. C. Santos, I. I. Bittencourt (Eds.), *Artificial Intelligence in Education*, volume 14829, Springer Nature Switzerland, Cham,

- 2024, pp. 324–337. URL: https://link.springer.com/10.1007/978-3-031-64302-6_23. doi:10.1007/978-3-031-64302-6_23, series Title: Lecture Notes in Computer Science.
- [6] P. Trotter, M. Vanderburg, R. Vanderburg, AI-Assisted Pedagogies: Enhancing Mathematical Literacy and Open-Ended Problem-Solving with ChatGPT, ASCILITE Publications (2024) 635–640. URL: <https://publications.ascilite.org/index.php/APUB/article/view/1443>. doi:10.14742/apubs.2024.1443.
- [7] A. Nitze, Future-proofing Education: A Prototype for Simulating Oral Examinations Using Large Language Models, 2024. URL: <https://arxiv.org/abs/2401.06160>. doi:10.48550/ARXIV.2401.06160, version Number: 1.
- [8] S. Amanda Puteri, Y. Saputri, Y. Kurniati, The Impact of Artificial Intelligence (AI) Technology on Students' Social Relations, BICC Proceedings 2 (2024) 153–158. URL: <https://biccproceedings.org/index.php/bicc/article/view/121>. doi:10.30983/bicc.v1i1.121.
- [9] V. J. Owan, K. B. Abang, D. O. Idika, E. O. Etta, B. A. Bassey, Exploring the potential of artificial intelligence tools in educational measurement and assessment, Eurasia Journal of Mathematics, Science and Technology Education 19 (2023) em2307. URL: <https://www.ejmste.com/article/exploring-the-potential-of-artificial-intelligence-tools-in-educational-measurement-and-assessment-13428>. doi:10.29333/ejmste/13428.
- [10] E. Kasneci, K. Sessler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günemann, E. Hüllermeier, S. Krusche, G. Kutyniok, T. Michaeli, C. Nerdel, J. Pfeffer, O. Poquet, M. Sailer, A. Schmidt, T. Seidel, M. Stadler, J. Weller, J. Kuhn, G. Kasneci, ChatGPT for good? On opportunities and challenges of large language models for education, Learning and Individual Differences 103 (2023) 102274. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1041608023000195>. doi:10.1016/j.lindif.2023.102274.
- [11] OpenAI, Europe privacy policy, 2024. URL: <https://web.archive.org/web/20250130122518/https://openai.com/policies/eu-privacy-policy/>.
- [12] L. Huang, Ethics of Artificial Intelligence in Education: Student Privacy and Data Protection, Science Insights Education Frontiers 16 (2023) 2577–2587. URL: <https://bonoi.org/index.php/sief/article/view/1084>. doi:10.15354/sief.23.re202.
- [13] Virginia Wesleyan University, Student Overview & Analytics in Learn Ultra, 2025. URL: <https://malboncenter.vwu.edu/services/academic-affairs-documentation/blackboard/blackboard-learn-ultra-documentation/student-overview-analytics-in-learn-ultra>, accessed: 2025-02-25.
- [14] UHasselt, Algoritmen en datastructuren (4900), 2023. URL: <https://studiegidswww.uhasselt.be/opleidingsonderdeel.aspx?a=2023&i=4900&n=4&t=01>.
- [15] Artificial Analysis, Llama 3.2 Instruct 3B: Intelligence, Performance & Price Analysis, 2025. URL: <https://artificialanalysis.ai/models/llama-3-2-instruct-3b>, accessed: 2025-02-04.
- [16] Gemma Team, A. Kamath, J. Ferret, S. Pathak, N. Vieillard, R. Merhej, S. Perrin, T. Matejovicova, A. Ramé, M. Rivière, L. Rouillard, T. Mesnard, G. Cideron, J.-b. Grill, S. Ramos, E. Yvinec, M. Casbon, E. Pot, I. Penchev, G. Liu, F. Visin, K. Kenealy, L. Beyer, X. Zhai, A. Tsitsulin, R. Busa-Fekete, A. Feng, N. Sachdeva, B. Coleman, Y. Gao, B. Mustafa, I. Barr, E. Parisotto, D. Tian, M. Eyal, C. Cherry, J.-T. Peter, D. Sinopalnikov, S. Bhupatiraju, R. Agarwal, M. Kazemi, D. Malkin, R. Kumar, D. Vilar, I. Brusilovsky, J. Luo, A. Steiner, A. Friesen, A. Sharma, A. Sharma, A. M. Gilady, A. Goedeckemeyer, A. Saade, A. Kolesnikov, A. Bendebury, A. Abdagic, A. Vadi, A. György, A. S. Pinto, A. Das, A. Bapna, A. Miech, A. Yang, A. Paterson, A. Shenoy, A. Chakrabarti, B. Piot, B. Wu, B. Shahriari, B. Petrini, C. Chen, C. L. Lan, C. A. Choquette-Choo, C. Carey, C. Brick, D. Deutsch, D. Eisenbud, D. Cattle, D. Cheng, D. Paparas, D. S. Sreepathihalli, D. Reid, D. Tran, D. Zelle, E. Noland, E. Huizenga, E. Kharitonov, F. Liu, G. Amirkhanyan, G. Cameron, H. Hashemi, H. Klimczak-Plucińska, H. Singh, H. Mehta, H. T. Lehari, H. Hazimeh, I. Ballantyne, I. Szpektor, I. Nardini, J. Pouget-Abadie, J. Chan, J. Stanton, J. Wieting, J. Lai, J. Orbay, J. Fernandez, J. Newlan, J.-y. Ji, J. Singh, K. Black, K. Yu, K. Hui, K. Vodrahalli, K. Greff, L. Qiu, M. Valentine, M. Coelho, M. Ritter, M. Hoffman, M. Watson, M. Chaturvedi, M. Moynihan, M. Ma, N. Babar, N. Noy, N. Byrd, N. Roy, N. Momchev, N. Chauhan, O. Bunyan, P. Botarda, P. Caron, P. K. Rubenstein, P. Culliton, P. Schmid, P. G. Sessa, P. Xu, P. Stanczyk, P. Tafti, R. Shivanna, R. Wu, R. Pan, R. Rokni, R. Willoughby, R. Vallu, R. Mullins,

- S. Jerome, S. Smoot, S. Girgin, S. Iqbal, S. Reddy, S. Sheth, S. Pöder, S. Bhatnagar, S. R. Panyam, S. Eiger, S. Zhang, T. Liu, T. Yacovone, T. Liechty, U. Kalra, U. Evci, V. Misra, V. Roseberry, V. Feinberg, V. Kolesnikov, W. Han, W. Kwon, X. Chen, Y. Chow, Y. Zhu, Z. Wei, Z. Egyed, V. Cotruta, M. Giang, P. Kirk, A. Rao, J. Lo, E. Moreira, L. G. Martins, O. Sanseviero, L. Gonzalez, Z. Gleicher, T. Warkentin, V. Mirrokni, E. Senter, E. Collins, J. Barral, Z. Ghahramani, R. Hadsell, Y. Matias, D. Sculley, S. Petrov, N. Fiedel, N. Shazeer, O. Vinyals, J. Dean, D. Hassabis, K. Kavukcuoglu, C. Farabet, E. Buchatskaya, J.-B. Alayrac, R. Anil, Dmitry, Lepikhin, S. Borgeaud, O. Bachem, A. Joulin, A. Andreev, C. Hardin, R. Dadashi, L. Hussenot, Gemma 3 Technical Report, 2025. URL: <https://arxiv.org/abs/2503.19786>. doi:10.48550/ARXIV.2503.19786, version Number: 1.
- [17] SBERT, Pretrained Models, 2025. URL: https://www.sbert.net/docs/sentence_transformer/pretrained_models.html, accessed: 2025-04-04.
- [18] M. Grootendorst, A. Mishra, A. Matsak, OysterMax, Priyanshul Govil, Y. Ogura, V. D. Warmerdam, Yusuke1997, MaartenGr/KeyBERT: v0.8, 2023. doi:10.5281/ZENODO.4461264.
- [19] J. Wei, X. Wang, D. Schuurmans, M. Bosma, b. ichter, F. Xia, E. Chi, Q. V. Le, D. Zhou, Chain-of-Thought Prompting Elicits Reasoning in Large Language Models, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), *Advances in Neural Information Processing Systems*, volume 35, Curran Associates, Inc., 2022, pp. 24824–24837. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf.