

# Performance of Large Language Models in a Computer Science Degree Program<sup>\*</sup>

Tim Krüger<sup>1</sup>[0009–0005–1876–8397] and Michael Gref<sup>1</sup>[0009–0003–9903–9959]

Niederrhein University of Applied Sciences, Krefeld, Germany  
tim.krueger@stud.hn.de  
michael.gref@hs-niederrhein.de

**Abstract.** Large language models such as ChatGPT-3.5 and GPT-4.0 are ubiquitous and dominate the current discourse. Their transformative capabilities have led to a paradigm shift in how we interact with and utilize (text-based) information. Each day, new possibilities to leverage the capabilities of these models emerge. This paper presents findings on the performance of different large language models in a university of applied sciences' undergraduate computer science degree program. Our primary objective is to assess the effectiveness of these models within the curriculum by employing them as educational aids. By prompting the models with lecture material, exercise tasks, and past exams, we aim to evaluate their proficiency across different computer science domains. We showcase the strong performance of current large language models while highlighting limitations and constraints within the context of such a degree program. We found that ChatGPT-3.5 averaged 79.9% of the total score in 10 tested modules, BingAI achieved 68.4%, and LLaMa, in the 65 billion parameter variant, 20%. Despite these convincing results, even GPT-4.0 would not pass the degree program - due to limitations in mathematical calculations.

## 1 Introduction

In the realm of natural language processing, large language models, hereafter only referenced as LLMs, have now become an integral part of our digital landscape. They have a widespread influence in today's discourse and a ubiquitous presence in various fields and industries [18]. Among these models, ChatGPT-3.5 and GPT-4.0 have emerged as prominent examples, captivating the attention of students, researchers, and developers alike. It is essential to look at these models in the context of higher education because they provide new ways and possibilities to teach, learn and perceive information. Useful for both students and instructors. They could help students, for example, by delivering a more personalized and interactive educational experience and acting as a kind of "learning buddy." For an instructor, the possibilities are also plenty. These models can generate supplementary materials, explanations, or examples [5].

---

<sup>\*</sup> Supported by funds of the *Bundesministerium für Bildung und Forschung*

Alternatively, they could aid in the assessment process by automating the grading procedure for all text-based requirements. A lot of research is currently taking place on this topic. For example, H. Gimpel et al. [5] have written an extensive essay on the opportunities but also the risks that generative AI models bring to higher education by collecting nearly 50 high-quality scholarly sources. They provide guidance for both students and instructors by providing hands-on recommendations for the usage of AI in higher education.

ChatGPT-3.5 and GPT-4.0 are not the only AI models impacting learning and teaching; much more software exists. [DeepL Write](#) can improve writing, from fixing grammar and punctuation mistakes to rephrasing entire sentences or sections. The same is true for [Grammarly](#), which offers users an AI text generation functionality for further improvements and suggestions regarding clarity, engagement, and delivery of a text. Even the creation of multimedia content is no problem. Programs like [Midjourney](#) and [Dall-E](#) allow users the creation of photorealistic images and visualizations with just a few prompts [10]. Furthermore, when Microsoft releases [Copilot](#), their AI support tool, with Office [11], the use of AI will have also arrived in all non-technical disciplines. These tools will then be used passively daily by millions of people, so we must look at the opportunities but also threats that these technologies can bring to higher education and learning/teaching in general.

As part of the research for this paper, we interviewed several professors from our teaching institution. We identified cheating and plagiarism as one of the main concerns. H. Gimpel et al. [5] have also dedicated several pages of their essay to this topic and stressed the importance of rules and guidelines that should be in place for the university environment without denying students access to this new technology. However, as the present work is limited to evaluating the performance of LLMs, this is an aspect to be explored in subsequent work. To accurately assess the benefits of this technology and their usage as educational aids, we set out to evaluate the performance across our undergraduate computer science curriculum. In total, we collected 40 data points, where one data point represents the performance of one LLM or LLM variant in one module of the degree program.

## 2 Related Work

In this section, we want to explore some of the various research efforts that have examined the performance of LLMs in the field of computer science. The results, some of which differ significantly, inspired us to test the performance of these LLMs in our degree program as well. Table 1 shows a small selection of test and exam results published by OpenAI [14], with the release of their GPT-4.0 model, and one exam (Algorithms and Data Structures) tested by Bordt et al. [2]. The first results are from LeetCode, a popular online platform that provides programming exercises and coding challenges commonly found in technical interviews [7]. The programming exercises and algorithmic problems are divided into three difficulty ranges (easy, medium, and hard) [13].

Table 1: Exam results for GPT-3.5/GPT-4.0 [14] [2].  
 Values are rounded to the first decimal place.

Test / Exam	GPT-3.5	GPT-4.0
LeetCode (Hard)	0.0%	6.6%
LeetCode (Easy)	29.3%	75.6%
Algorithms & Data Structures	51.3%	60.0%

The platform is aimed at software developers and programmers to enhance their programming skills by solving algorithmic problems [21].

In the *easy* problem section, GPT-3.5 answered 12 out of 41 questions correctly, resulting in a performance of 29.3%. GPT-4.0 answered 31 out of 41 questions correctly, resulting in a performance of 75.6% - an improvement of 47.3 percentage points [14]. The results may depend on the exact category and programming language [13]. Nikolaidis et al. found that in their case, ChatGPT-3.5 solved 45% of 50 randomly selected *easy* LeetCode problems correctly while providing noticeably better results in the programming languages Java and Python.

When tested by OpenAI, GPT-3.5 could not solve a single of the *hard* problems on LeetCode [14]. These results again may depend on the type of problem that had to be solved [13]. Nikolaidis et al. found that ChatGPT-3.5 solved 10 out of 21 *hard* problems correctly, resulting in 47.6% accuracy. ChatGPT-3.5 would then, in fact, even severely outperform GPT-4.0 when tested by OpenAI, which was able to solve 3 out of 45 *hard* problems correctly (6.6% accuracy) [14].

Bordt et al. tested ChatGPT-3.5 and GPT-4.0 on an undergraduate computer science exam in *Algorithms and Data Structures*. The exam was fed to the LLMs in the same way students would receive it. The answers of the models were transferred to paper by the testers and mixed with the solutions of the students [2]. ChatGPT-3.5 scored 20.5 out of 40 possible points (51.25%), allowing it to pass the exam narrowly. GPT-4.0 improved that score by 8.75 percentage points, reaching 60% (24/40 points). With this result, GPT-4.0 outperforms the average student, which scores 23.9 in the mean [2].

Both ChatGPT-3.5 and GPT-4.0 indicate wide-ranging capabilities in the field of computer science. GPT-4.0 also seems to be an improvement over GPT-3.5 in every way. The findings on performance variation are worth noting for our research. The LLMs' answers seem to depend on the corresponding computer science discipline category and the specific programming language asked [13]. It is also relevant to note that the programming errors generated seem to be mainly semantic. The models hardly make syntax errors, but the code, if wrong, can have serious logic errors [13].

### 3 Methodology

The crux of our methodology is the evaluation of various LLMs by feeding them academic content drawn from a bachelor's degree program in computer science at a university of applied sciences.

Table 2: Overview of our exam data set.

Semester	Written exam	Oral exam
1	0	0
2	2	0
3	2	0
4	2	1
5	2	1
Sum	8	2

We aim to determine each model’s overall performance and identify the highest and lowest-scoring modules, grade distributions, and potential affinities for certain topics. Additionally, the study aims to determine whether the models would complete the degree program. Our data set comprised samples of past exams from ten different modules of the degree program, see Table 2. This core data set was complemented with information from questionnaires, practice exercises, and lecture notes to offer a more holistic view of the curriculum. For modules with oral exams, the questions were based on the same data but created in consultation with the supervising professor to simulate realistic exam scenarios. Only examinations for which the professors gave their approval were taken into account for the study.

The criteria for evaluation were adapted for each module. In written exams, we employed the evaluation system and point allocation provided by the supervising professor. In oral exams, we weighted questions according to complexity and difficulty. These questions were finalized in consultation with the supervising professors. Evaluating and assessing the performance involved verifying correctness, compiling and testing program code, and recalculating solutions.

Due to the limitations of certain models in handling multimedia input, we partly excluded those tasks from our assessment. Adjustments were made to the total score and weighting of the exam accordingly. In instances where it was feasible, we transformed such tasks into a suitable textual format with tables and data structures being converted to markdown and diagrams re-imagined into the corresponding UML representation. We tested ChatGPT-3.5, AI-powered Bing (referred to as BingAI) [12], StableLM-Alpha in the 7 billion parameter version, and LLaMa in both the 7 billion and 65 billion parameter versions. Towards the end of our project, we also received access to GPT-4.0 but were restricted in using this model due to time constraints. We viewed these selections as an appropriate mix of open- and closed-access LLMs. StableLM includes various LLMs published by Stability AI. The size of these models ranges from 3 billion to 65 billion parameters. A 175 billion parameter variant is also planned [16]. The models are published in different versions and trained on different datasets. We use the StableLM-Alpha-7B variant, which was trained on a dataset based on *The Pile* [3]. All models are hosted on *The Hugging Face Hub*, and some are accessible through a web interface [17].

LLaMa refers to a collection of different LLMs ranging from 7B to 65B parameters, published by MetaAI [19]. We used LLaMa with the project *llama.cpp*, an open-source C/C++ port of several LLMs [4]. This project supports 8-bit, 5-bit, and 4-bit integer quantization, a technique that significantly reduces the memory requirements of the models. In the case of LLaMa, this allowed us to run the models in RAM instead of GPU memory. We considered this approach a more realistic simulation, as the models otherwise require a significant amount of GPU memory. However, there is the possibility of a degradation in model accuracy. There seems to be a trade-off between model size and quality, depending on the quantization method [22]. In the case of LLaMa-7B, the file size got reduced from 13 GB when using 16-bit floats to 3.5 GB when using 4-bit integer quantization. The perplexity [6] rose from 5.9066 to 6.1565, an increase of only 4.23% [4].

The prompting of the models was a carefully considered aspect of this research project. Prompt engineering has been shown to improve the performance of models in various studies (e.g. [20]). However, to provide a broad overview of the performance across the curriculum, we opted to prompt all models only once and use the first response provided by each model. Before starting the assessment, a generic pre-prompt was given in each case, setting the context that they were interacting in a simulated exam scenario and outlining expectations for responses.

*I am now going to ask you a few questions from a hypothetical [insert topic or subject] exam of an undergraduate computer science degree program. I want you to answer the questions to the best of your knowledge and capabilities. Please answer briefly and concisely unless I explicitly ask for a more detailed answer! Please answer purely in continuous text or bullet points. If output in chart or table form is desired, I will let you know.*

## 4 Experimental Results

We tested, between May and July of 2023, ten modules each with ChatGPT-3.5 and BingAI, six modules with GPT-4.0, and fourteen modules in total with StableLM-Alpha-7B, LLaMa-7B, and LLaMa-65B, resulting in forty data points. We have not been able to test every model iteration on every module of the curriculum due to the time constraints of this project. All exam questions were adopted without modification. The following data underpins what we present as a comprehensive insight into the performance of these models across an array of computer science curriculum modules.

Referring to grades in the following, we calculated them according to the *modified Bavarian formula* corresponding to the German grading system [15]. Depending on the university, a conversion may be necessary. If not stated otherwise, 50% of the score are required to pass the exam.




Fig. 1: Exam results for Operating Systems (OS) and Object Oriented Application Development (OOA)

#### 4.1 1st Semester

We have neither received full approval nor the required content for any of the modules of the first semester from the responsible professors. This is left to be explored in subsequent work.

#### 4.2 2nd Semester

We have received approval for two second-semester modules, Operating Systems, and Object-Oriented Application Development. Figure 1 shows the exam results for each LLM in these modules. Operating Systems (OS) is a five credit-point module. In the module, students learn the structure of a modern operating system and algorithms and strategies for managing and allocating resources in it. They also develop programs in a UNIX environment and work out solutions to problems of interprocess communication [8]. Object Oriented Application Development (OOA) is a seven credit-point module. It focuses on teaching the methods and techniques of object-oriented programming.

Requirements are implemented using efficient algorithms and data structures. Programming is done in C++ [8].

ChatGPT-3.5 and BingAI performed quite well in OS, scoring 82.4% and 80.6%, respectively. The 7B parameter models performed significantly worse. LLaMa-7B-Q (quantized) scored 21.8%, and StableLM-7B scored 9.4%. While StableLM-7B could answer almost no questions, LLaMa-7B-Q could still answer questions about shell commands and general operating system terms. Nevertheless, it was too little to pass the exam. ChatGPT-3.5 and BingAI were able to answer many questions. The models did make mistakes when calculating memory usage and applying paging algorithms. This cost them points but kept the good result the same. ChatGPT-3.5 passed this exam with a grade of 2.0 and BingAI with a Grade of 2.1. OOA is the only module in which BingAI performed better than ChatGPT-3.5 in all our testing. The latter LLM scored 75%, whereas BingAI scored 86.3%. This results in grades of 2.5 for ChatGPT-3.5 and 1.8 for BingAI. The score for BingAI is one of the best results for this LLM in all our tests. The biggest problems ChatGPT-3.5 had were with implementing the object-oriented interfaces in C++. The code compiled but either didn't work as it should or implemented something completely different from the task. StableLM-7B performed slightly better than LLaMa-7B-Q in this exam. However, both LLMs had severe problems with the assignments, solving almost no tasks.

### 4.3 3rd Semester

We have received approval for two modules of the third semester, Web Engineering and Distributed Systems. The results for the modules of the third semester can be seen in Figure 2. Web Engineering (WEB) is a five credit-point module. It covers the technical fundamentals of modern web-based technologies and architectural, development, and analysis tools for web-based systems. On the front end, students in this module work with HTML, CSS, and JavaScript; on the backend side, with a mixture of Javascript and Python [8]. A student with 33% or more would pass the exam, as determined by the supervising professor. Distributed Systems (DS) is also a five credit-point module. Students of this module learn about distributed system architectures and techniques for synchronization and communication. At the end of this module, they can design, implement and evaluate their own distributed computing structures. The implementation within this module takes place in C/C++ [8]. Our set exam consists of questionnaire material.

WEB was one of the best exam results in all our testing for ChatGPT-3.5, scoring an even 1.0 on the exam with 98.3%. Even the most extensive task, a partial Python implementation of a backend server for the membership management of a business, was solved completely and correctly. BingAI was also able to answer almost every question correctly. Only in the implementation part did BingAI make logical errors and omit required functionalities. This still resulted in 90% or a grade of 1.4. StableLM-7B and LLaMa-7B-Q had surprisingly massive problems in this exam, despite the extensive question part.




Fig. 2: Exam results for Web-Engineering (WEB) and Distributed Systems (DS)

Almost no question could be answered completely or correctly. Both models also failed the implementation part. LLaMa-7B-Q scored 15.8%, slightly better than StableLM-7B, with 9.2%. In DS, both ChatGPT-3.5 and BingAI performed worse than in WEB. ChatGPT-3.5 got a grade of 2.2 with a result of 78.5%, and BingAI a 2.8 with 70%. The models could answer almost all simple or introductory questions to the topic correctly but had problems with more in-depth questions, e.g., on network data formats or broker implementations. DS was the first module in which we tested GPT-4.0. With a result of 95.5%, it got a grade of 1.2 and topped the grade of ChatGPT-3.5 by a whole level. GPT-4.0 answered almost every question completely and correctly in this exam. StableLM-7B scored 30% in DS, the best result for this LLM in all our tests. Surprisingly, it was able to answer difficult questions on *CORBA*, *SOAP* interfaces, and synchronization mechanisms but failed on simpler, more general questions, like resilience and fault tolerance of distributed systems. Otherwise, it would have had a real chance to pass the exam. In this module, we also tested LLaMa-65B-Q for the first time. With a result of 14.5%, it performed only slightly better than LLaMa-7B-Q with 13.5%. Considering the difference in size, this is a disappointing result.



#### 4.4 4th Semester

We have received approval for three modules of the fourth semester: Data Network Management, Interactive Systems, and Numerical Analysis. The results can be taken from Figure 3. Data Network Management (DNM) is a six credit-point module. It provides in-depth, application-oriented knowledge of network administration. Students in this module acquire skills in the design, development, and deployment of large-scale computer networks, as well as techniques for securing them [8]. Interactive Systems (IAS) is a five credit-point module. It focuses on software ergonomics and the design and implementation of portable interactive systems. Students of this module learn how to model application-oriented and ergonomic human-machine interfaces [8]. The implementations in this module are web-based in the programming languages JavaScript and Python. As determined by the supervising professor, the module is considered to be passed if 33% of the total points are achieved. Numeric Analysis (NUM) is an elective course in our computer science bachelor's degree program, which gives five credit points. Topics covered include computer arithmetic and rounding errors, systems of linear equations, and linear equilibrium calculus. The module is concluded with an oral examination [8]. This is one of the modules in which we simulated an exam by taking questions from a questionnaire.

DNM is the first module in our tests in which even the larger LLMs have experienced problems. ChatGPT-3.5 barely passed the exam with 51.9% or a grade of 3.8. BingAI had even more difficulties and failed the exam with a score of only 47.1%. The application of firewall rules and routing protocols for custom multi-area networks presented in the exam was particularly problematic for both LLMs. GPT-4.0 performed the best in this exam. It was also unable to completely solve the more difficult tasks but often provided correct partial solutions or made less serious errors than the other two LLMs. GPT-4.0 passed the exam with 65.4% or a grade of 3.0. IAS went very well for ChatGPT-3.5. With 96.7%, it got a grade of 1.1. It answered almost all comprehension and knowledge questions correctly. Even more complex tasks, such as the design of a user interface, were solved completely and correctly. BingAI performed almost 30 percentage points worse in this exam, resulting in one of the biggest gaps between these two LLMs in all our testing. It got a grade of 2.4, or 68.8% of the total score. The grade of 2.4 comes from the fact that the exam is considered passed from 35%, and larger results are offset by the formula linearly. Nevertheless, BingAI had problems with several questions in this exam and either answered incorrectly or omitted information. StableLM-7B and LLaMa-7B-Q had no chance of passing this exam, with a performance of 6.7% and 13.3%, respectively. Nearly every answer had massive errors or large information gaps. The models also lost context in between and started talking about completely different topics. In the simulated oral exam on numerical analysis, mainly comprehension questions were asked, and hardly any calculations had to be done. This led to excellent results for both GPT models. ChatGPT-3.5 got a grade of 1.6, with 90% of the total score, whereas GPT-4.0 increased this to 95% and a grade of 1.3.




Fig. 3: Exam results for Data Network Management (DNM), Interactive Systems (IAS) and Numerical Analysis (NUM)

Both models could answer almost every question completely and correctly and only made minimal errors. BingAI had great problems in this exam, although it was mainly about knowledge reproduction, and scored well behind the GPT models with 68% of the total score, or a grade of 2.9. BingAI had problems with several questions and made mistakes while reproducing information. For example, when asked about the complexity of the *Gauss Algorithm*, BingAI gave a reference to Wikipedia but then misquoted the article with a complexity of  $O(n^2)$ .

#### 4.5 5th Semester

We have received approval for three modules of the fifth semester: Data Science, Software Engineering, and Real-Time Systems. The results can be seen in Figure 4. Data Science (DSC) is an elective course with five credit points. The module provides an introduction to Big Data and Machine Learning. Students of this module will learn to extract, prepare and analyze large data sets [8].




Fig. 4: Exam results for Data Science (DSC), Software Engineering (SWE) and Real-Time Systems (RTS)

The module concludes with an oral exam which we simulated by taking questions from a questionnaire. Software Engineering (SWE) is a five credit-point module. It covers advanced solutions for building, testing, and maintaining large IT systems and techniques for organizing big software projects. A special focus is on effort estimation and (agile) software development processes [8]. As determined by the supervising professor, the module is considered to be passed if 33% of the total points are achieved. Real-Time Systems (RTS) is a five credit-point module. It focuses on the architecture, the concepts, and the functionalities of modern real-time systems. Students learn aspects of concurrent real-time programming and how to deal with time constraints and task management [8]. The module is concluded with a written exam in which a special focus is placed on manual real-time proof for various scheduling methods.

DSC is the best-performing module for ChatGPT-3.5 in all our tests. With 99%, ChatGPT-3.5 scored a 1.0. Every question was answered correctly, from problems in the field of *Big Data* to data preparation, to classification and clustering methods. DSC is also the only module in which ChatGPT-3.5 outperformed GPT-4.0.

The latter model did not score itself badly with 94% and a score of 1.3, but unfortunately gave partially wrong answers to questions about *Eventual Consistency* and *Sharding*. Such results are possible since we only prompt all models once in our tests. BingAI performed again worse than the GPT models. With 81%, it achieved a score of 2.1. It was challenging for BingAI to make its own decisions in tasks, e.g., choosing between an aggregate-oriented or a relational data model. In SWE, the GPT models were again well ahead of BingAI. ChatGPT-3.5 achieved a 1.9 with 78.1%, GPT-4.0 a 1.7 with 82.9%, while BingAI only achieved a 2.9 with 56.7%. The GPT models answered most of the questions completely and correctly but made massive errors in designing test cases for a finite state machine. BingAI could not solve this task either and made errors in explaining design patterns and performing an effort estimation using *Function Point Analysis*. StableLM-7B and LLaMa-7B-Q performed poorly in this module, scoring 3.8% and 6.2%, respectively. LLaMa-65B-Q performed significantly better than the 7 billion parameter version and nearly passed the exam with 25.2%, out of 33% needed. Interestingly, LLaMa-65B-Q was able to partially answer difficult questions on software development principles and the design of component diagrams but failed to explain unit tests. RTS is a demanding exam in which many calculations have to be done. Computational time requirements, core workloads, and a large, manual real-time proof must be calculated. This exam has proven to be extremely difficult for all models tested. ChatGPT-3.5 scored 29.4%, GPT-4.0 scored 33.8%, and BingAI scored 23.5%. Accordingly, all models failed the exam. The models could answer a few simple introductory questions but, early on, miscalculated the computational kernel allocation for a round-robin scheduling procedure. No model was able to solve this task correctly. Likewise, no model was able to calculate the real-time proof correctly. This task is nested, with each intermediate calculation evaluated individually but often a prerequisite for the next calculation. BingAI lost all context in this task after the third partial calculation. The GPT models could continue to calculate but miscalculated fatally early on. Both the calculation path and the result were not correct. When calculating the average execution time for processes of a machine, both ChatGPT-3.5 and GPT-4.0 set up the correct formula, adding all times and dividing by the amount, but only GPT-4.0 also got the correct result; ChatGPT-3.5 miscalculated. All in all, the models were heavily overcharged with this exam.

## 5 Discussion

An overview of the average performances and results achieved across all tested modules can be seen in Table 3. It is important to note that all data collected is only a snapshot that considers the systems' performance at the time of the assessment. These systems continue to evolve. ChatGPT-3.5 achieved an average of 79.9% of the maximum possible score in the ten modules tested. ChatGPT-3.5 performed particularly well in modules with a high proportion of web development or high-level programming language content, like Python and JavaScript.

Table 3: Summary of the average performance proportionally calculated to all modules taken.

Model	Average score	#Modules	Passed/Failed
GPT-4.0	80.2%	6	5/1
ChatGPT-3.5	79.9%	10	9/1
BingAI	68.4%	10	8/2
LLaMa-65B-Q	20.0%	2	0/2
LLaMa-7B-Q	12.3%	6	0/6
StableLM-7B	10.8%	6	0/6

Even in the field of data science, ChatGPT-3.5 achieved almost full marks. In exams with more complex tasks like Operating Systems or Data Network Management, ChatGPT-3.5 often provided at least an approach to the solution. We noted major difficulties for this LLM with various tasks that required mathematical calculations. The application of scheduling algorithms and the calculation of core utilization and process runtimes posed significant challenges for ChatGPT-3.5. Due to these shortcomings, passing the Real-Time Systems exam is not currently possible. Accordingly, the LLM would not be capable of completely finishing our bachelor’s degree program in computer science. However, with an understanding of its strengths and weaknesses, ChatGPT-3.5 shows great potential as an outstanding learning aid for students and lecturers.

GPT-4.0 achieved even better results than ChatGPT-3.5, obtaining an average performance of 80.2%. This score is expected to increase even further if the missing modules are tested with GPT-4.0. A strong focus on specific programming languages or fields of computer science, as observed in ChatGPT-3.5, could not be detected in GPT-4.0. At the same time, GPT-4.0 demonstrated a more consistent overall performance. Like ChatGPT-3.5, GPT-4.0 had difficulties with tasks that required calculations; this also resulted in the failure of the Real-Time Systems exams. For the same reason, GPT-4.0 would not be able to finish the degree program. However, it should be noted that GPT-4.0, despite the identified difficulties, represents an improvement in all areas over ChatGPT-3.5. The use of plugins, for instance, to redirect mathematical computations to a system like WolframAlpha could significantly improve this outcome. This is left to be explored in future studies.

BingAI scored much lower than the GPT models in our tests, with 68.4%. It was the only one of these three models that failed two exams rather than just one, with one of the exams (Data Network Management) being not calculation-intensive. BingAI often encountered problems when the solutions were not directly searchable online. Thus, it made mistakes in extracting information from texts or creating and presenting solutions. Even when the answer to a question could be found via an internet search, BingAI sometimes made inexplicable citation errors. BingAI also provided the shortest responses of all the LLM systems tested, often ignoring aspects of a question. According to the current state, BingAI is inferior to the GPT models.

LLaMa-7B-Q showed poor results, with an average performance of 12.3% in six tested modules. It often had difficulties understanding questions, lost context, or started talking about completely different topics. LLaMa-7B-Q could not solve a single task of an exam. According to our tests, it would not be possible for this LLM to pass any module.

LLaMa-65B-Q showed better results with an average performance of 20.0%, but was also tested only in two modules. It scored one percentage point and 19 percentage points better than its 7 billion parameter counterpart. At this point, more tests are needed to make a final statement about the performance differences between these models. Nevertheless, a trend can be determined: LLaMa-65B-Q performs significantly worse than BingAI, let alone the GPT models. After our tests, whether it would pass a single module is questionable, and it is not suitable for use as a learning aid.

StableLM-7B, tested in six modules, achieved the worst results of all tested LLMs with 10.8%. It was unable to answer any question correctly and completely. Interestingly, StableLM-7B often related questions to a business context or attempted to answer them in such a context. StableLM-7B even understood complex questions from the field of project management but could not establish a reference to computer science or software development. According to our tests, this LLM is also unsuitable as a learning aid.

## 6 Conclusion

In the presented study, we tested and evaluated the performance of various LLMs across a series of modules in a bachelor’s computer science degree program. Our results are in line with existing research (e.g. [1]) by showing strong performances of Generative Pre-training models (GPT) across an undergraduate curriculum while having severe restrictions in key areas. A prevalent worry is the potential for essays to progressively lose significance as evaluation tools within higher education [5]. Our tests show the strength and topic affinities of current LLMs, but also their weaknesses, distinctively in mathematical computations. We conclude from these results that a comprehensive blueprint for our curriculum remains elusive at this point. Despite this, the deployment of these models presents lecturers with challenges, as the detection of plagiarism in AI-generated content is not particularly mature yet [9]. It is imperative to recognize that the GPT models in our tests have completed numerous examinations with scores above 95%. Given that some of our examination rules allow aids, and the pattern of past exams often remains unchanged, the sophisticated capabilities of these models could potentially create near-perfect and legally permissible "cheat" sheets. This, combined with the advancing abilities of current LLMs [5], compels us to reconsider and construct robust examination methods. Oral and written exams without aids remain valid alternative options [5].

The smaller models in our tests exhibit substantial performance deficiencies, with profound disparities encountered in almost all performance-defining areas. Consequently, they currently do not measure up as viable educational aids.

Continued research may examine the performance of existing models in unexplored curriculum modules. Furthermore, additional modules could be examined to provide a broader overview. Future research could also extend to the study of other LLMs, such as [Google Bard](#). Also, broadening the scope to related disciplines, like electrical engineering, would be beneficial to gain a better understanding of domain-specific performance capabilities.

## References

1. Binz, M., Schulz, E.: Using cognitive psychology to understand GPT-3. Proceedings of the National Academy of Sciences (2023)
2. Bordt, S., von Luxburg, U.: ChatGPT Participates in a Computer Science Exam (2023)
3. Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., Presser, S., Leahy, C.: The Pile: An 800gb dataset of diverse text for language modeling. arXiv preprint arXiv:2101.00027 (2020)
4. Georgi Gerganov: llama.cpp Github Repository. <https://github.com/ggerganov/llama.cpp> (2023), [Online; accessed 03-July-2023]
5. Gimpel, H., Hall, K., Decker, S., Eymann, T., Lämmermann, L., Mädche, A., Röglinger, M., Ruiner, C., Schoch, M., Schoop, M., Urbach, N., Vandrik, S.: Unlocking the power of generative AI models and systems such as GPT-4 and ChatGPT for higher education: A guide for students and lecturers. Hohenheim Discussion Papers in Business, Economics and Social Sciences 02-2023, Universität Hohenheim, Fakultät Wirtschafts- und Sozialwissenschaften, Stuttgart (2023), <http://hdl.handle.net/10419/270970>
6. Gonen, H., Iyer, S., Blevins, T., Smith, N.A., Zettlemoyer, L.: Demystifying Prompts in Language Models via Perplexity Estimation (2022)
7. Harper, J.: Interview insight: How to get the job. In: A Software Engineer's Guide to Seniority: A Guide to Technical Leadership, pp. 19–28. Springer (2022)
8. Hochschule Niederrhein: Modulhandbuch zum Vollzeit Studiengang Bachelor Informatik nach Prüfungsordnung 2013. [https://www.hs-niederrhein.de/fileadmin/dateien/FB03/Studierende/Bachelor-Studiengaenge/P02013/modul\\_\\_bi.pdf](https://www.hs-niederrhein.de/fileadmin/dateien/FB03/Studierende/Bachelor-Studiengaenge/P02013/modul__bi.pdf) (2019), [Online; accessed 26-June-2023]
9. Khalil, M., Er, E.: Will ChatGPT get you caught? Rethinking of Plagiarism Detection (2023)
10. Krüger, N.: Künstliche Intelligenz in Training, Weiterbildung und Beratung: 70 direkt anwendbare und erprobte KI-Tools. pitchnext GmbH via Kindle Direct Publishing (2022)
11. Microsoft: Introducing Microsoft 365 Copilot — A whole new way to work. <https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot-a-whole-new-way-to-work/> (2023), [Online; accessed 03-July-2023]
12. Microsoft: Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web. <https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web> (2023), [Online; accessed 04-September-2023]
13. Nikolaidis, N., Flamos, K., Feitosa, D., Chatzigeorgiou, A., Ampatzoglou, A.: The End of an Era: Can Ai Subsume Software Developers? Evaluating Chatgpt and Copilot Capabilities Using Leetcode Problems. Evaluating Chatgpt and Copilot Capabilities Using Leetcode Problems (Preprint) (2023)

14. OpenAI: GPT-4 Technical Report. <https://arxiv.org/pdf/2303.08774.pdf> (2023)
15. Sekretariat der ständigen Konferenz der Kultusminister der Länder in der Bundesrepublik Deutschland: Vereinbarung über die festsetzung der gesamtnote bei ausländischen hochschulzugangszugnissen. [https://www.kmk.org/fileadmin/Datien/pdf/ZAB/Hochschulzugang\\_Beschluesse\\_der\\_KMK/GesNot05.pdf](https://www.kmk.org/fileadmin/Datien/pdf/ZAB/Hochschulzugang_Beschluesse_der_KMK/GesNot05.pdf) (2013), [Online; accessed 24-June-2023]
16. Stability AI: StableLM Github Repository. <https://github.com/Stability-AI/StableLM> (2023), [Online; accessed 03-July-2023]
17. Stability AI: StableLM hugging face. <https://huggingface.co/stabilityai> (2023), [Online; accessed 03-July-2023]
18. Statista: Generative AI - Worldwide. <https://www.statista.com/outlook/tmo/artificial-intelligence/generative-ai/worldwide> (2023), [Online; accessed 05-July-2023]
19. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., Lample, G.: LLaMA: Open and Efficient Foundation Language Models (2023)
20. White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., Schmidt, D.C.: A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT (2023)
21. Wikipedia contributors: Leetcode — Wikipedia, the free encyclopedia. <https://en.wikipedia.org/w/index.php?title=Leetcode&oldid=1159537770> (2023), [Online; accessed 19-June-2023]
22. Yao, Z., Wu, X., Li, C., Youn, S., He, Y.: ZeroQuant-V2: Exploring Post-training Quantization in LLMs from Comprehensive Study to Low Rank Compensation (2023)