Preprints.org

ChatGPT: The End of Online Exam Integrity?

Teo Susnjak *

Article

ChatGPT: The End of Online Exam Integrity?*

Teo Susnjak 🆔

School of Mathematical and Computational Sciences, Massey University, Albany, New Zealand; t.susnjak@massey.ac.nz

Abstract: Background: This study addresses the burgeoning impact of Large Language Models (LLMs) such as ChatGPT on the integrity of online examinations and seeks to demonstrate latent and advanced LLM reasoning capabilities. The research devises an iterative self-reflective strategy for invoking critical thinking and higher-order reasoning in LLMs based on multimodal exam questions using visuals and text. Methods: The proposed strategy was demonstrated and evaluated on real exam questions by subject experts and the performance of ChatGPT (GPT-4) with vision was estimated on an additional dataset of 600 text descriptions of multimodal exam questions, focusing on its response to complex, visual, and textual prompts. Results: The results indicate that latent multi-hop reasoning capabilities can be invoked using the proposed self-reflective strategy, which can effectively steer ChatGPT towards correct answers by integrating critical thinking from each modality into the final response. Meanwhile, ChatGPT demonstrated considerable proficiency in being able to answer multimodal exam questions across 12 subjects. Conclusions: This work challenges findings from recent studies which claim that existing LLMs are unable to reason on complex multimodal tasks, and therefore this study underscores the need for enhanced online exam security measures like proctoring software, that can mitigate the potential misuse of AI technologies in educational settings.

Keywords: AI impact on education; ChatGPT in exam cheating; multimodal online assessments; GPT-4 vision evaluation, educational integrity; Generative AI and examination security; large language models and assessments; critical thinking in LLMs; multi-hop reasoning in LLMs

---

1. Introduction

The landscape of higher education has experienced a significant transformation towards online learning, a trend that the recent pandemic has notably accelerated [1]. Higher Education Institutions (HEIs) worldwide have swiftly adapted to this new norm, transitioning to online classes and examinations to overcome these challenges [2–4]. This shift is expected to endure, driven by the recognized advantages of remote learning for both educational institutions and students [1,5].

With the increased adoption of online education, concerns regarding academic integrity have intensified [1,6,7]. The move to online assessments, in particular, has spotlighted the potential for increased cheating and academic misconduct [4,8–11], fueled by factors like anonymity, reduced supervision, and easier access to unauthorized resources during exams. Cheating in HEIs is a long-standing issue and has been compounded by the online modality, with faculty facing significant barriers to effectively countering dishonest practices [12,13]. Empirical research indicates a notable rise in academic dishonesty within online settings. For instance, Malik *et al.* [14] observed a sudden and significant increase in academic performance among students during the pandemic, attributing it to the facilitation of cheating in online exams, while Newton and Essex [15] noted recently that cheating in online exams is prevalent, emphasizing the allure of opportunity for dishonesty in these settings. Historical data corroborate the extensive nature of academic cheating, with studies revealing a significant proportion of students admitting to such behaviour [16,17].

These concerns around negative trends in academic integrity in online exams were already well advanced when ChatGPT was released, and a new and significant compounding factor was introduced.

---

*This paper is an updated version of the original paper published in December 2022, found here: https://arxiv.org/abs/2212.09292

ChatGPT's ability to answer exam questions has already been demonstrated in multiple studies [18]. The opportunity to cheat is now magnified making it easier than ever before, while the current AI technologies are more effective at enabling this than anything else beforehand. Richards *et al.* [19] showed that it is not even necessary for students to augment ChatGPT responses with their own work or edit its outputs to reach at least a passing grade, while students who augment the responses with their own additional material are highly likely to score even higher on their assessments. Newton and Essex [15] confirm what is self-evident, that students appear to be most likely to cheat in online exams when there is an opportunity to do so. In spite of this, studies [20] continue to claim that even unproctored online exams can still provide meaningful assessments of student learning. Much has been made of limitations of Large Language Models (LLMs) like ChatGPT due to their proclivity to fabricate non-existent facts or inappropriate information [18,21–23] referred to as hallucination, and with this comes the temptation to downplay and dismiss the performance capabilities of LLMs in exam contexts. It is also true that LLMs still encounter persistent difficulties across complex reasoning tasks [24] from which it argued and claimed that due to this, they are again unsuitable for answering sophisticated exam questions with reliability [25,26]. While it is therefore possible to reduce the threats posed by LLMs due to these technical grounds and limitations, some studies altogether dismiss the need for additional mitigation strategies that attempt to prevent cheating, based on principle [27] and instead shift the blame on to the faculty and other externalities instead, which effectively absolves students from any personal responsibility.

This matter is clearly of some contention; however, there are numerous studies that do take the urgency to devise additional and more effective countermeasures in this new context seriously. Studies have discussed mitigations like the use of LLM text-detection tools [28]. In an effort to uphold academic integrity HEIs have implemented various measures, including technological interventions like digital proctoring systems [6,29,30]. Other studies have called for recasting examination and assessment questions so that they require higher-order thinking and critical reasoning skills [31,32] instead of merely factual recall, others [33–35] have recommended integrating multimodal approaches into exams which venture beyond text-based and combine also visual interpretation tasks. Indeed, very recent studies [32,36] have demonstrated that LLMs perform poorly on complex reasoning tasks that combine both visual and text-based modalities into exam questions.

Aims and Contribution

This paper aims to show that LLMs are in fact more capable of performing complex reasoning tasks on multiple modalities than is currently believed. This work proposes a novel multimodal, iterative self-reflective strategy that decomposes a complex reasoning task into sub-tasks which are applied to each modality first, in order to guide the LLM towards a correct response. Therefore, this study empirically demonstrates how one of the last bastions of evading cheating via LLMs which involved constructing multimodal exam questions, is now also compromised. This study also extensively probes the visual understanding capabilities of state-of-the-art LLMs to answer multimodal exam questions across multiple subjects and disciplines seeking to probe which types of questions pose the most difficulties for the current LLMs. This paper also aims to critically examine the arguments and the inadequacies of current mitigations against cheating on online exams via LLMs, like detection tools for AI-generated text, while challenging the assumption that LLMs' reasoning limitations preclude effective cheating that is sufficient to pass various exams, finally concluding with a list of up-to-date recommendations on conducting online exams.

2. Background

The literature review addresses the challenges in distinguishing AI-generated text from human-authored content with the evolving capabilities of LLMs which are increasingly able to bypass advanced detection methods, focusing on the implications for academic integrity in online exams. The review also touches on LLM hallucination, noting that despite their propensity to generate misleading information,

they nonetheless often meet minimal academic assessment criteria, thus their occasional fallacious responses are not a barrier to them being used for cheating. Additionally, this section evaluates debates around LLMs' reasoning abilities, acknowledging criticisms of their limitations in complex reasoning tasks while recognizing emerging evidence of their enhanced capabilities through sophisticated, multi-hop reasoning strategies that need to be investigated on multimodal tasks.

2.1. Challenges in Detecting AI-Generated Text

Discerning AI-generated text from human-authored content continues to present formidable challenges; this is particularly true as LLMs attain unprecedented levels of stylistic fluency. The effectiveness of contemporary detection strategies is frequently undermined by the evolving sophistication of generative models, which have been shown to elude both well-established and nascent detection technologies with increasing ease [37]. Traditional methods, such as stylometric analysis, which were once reliable, now often fail to detect the subtle inconsistencies that are hallmarks of AI-generated texts; these inconsistencies include but are not limited to, atypical semantic patterns, unusual word choices, and subtle logical inconsistencies, thereby demonstrating the models' ability to produce text with high grammatical correctness and contextual appropriateness. The landscape is further complicated by adversarial training techniques, which allow LLMs to adapt and learn patterns that specifically avoid detection, thus challenging the efficacy of even the most advanced machine learning-based detectors [37]; these systems grapple with both false positives—mistakenly flagging human-written content—and false negatives—failing to detect AI-generated outputs [38]. This ongoing adaptation among LLMs initiates a cycle of challenge and response that characterizes the arms race between AI technologies and detection methodologies [39]. Thus, AI text detection techniques that are fine-tuned for one generation of LLMs are likely to become obsolete as new models emerge and proliferate.

2.2. LLM Hallucination Is Not an Insurmountable Problem

LLMs exhibit a remarkable capacity to produce fluent, coherent, and persuasive text, often accompanied by a tendency for it to be factually incorrect, misleading, or nonsensical information otherwise referred to as hallucination. Hallucination in LLMs should not come as a surprise - LLMs hallucinate by design [22,40]. Indeed, studies [18,25,32,41,42] have identified this as a major obstacle in using LLMs for producing correct and reliable responses in various critical contexts, and this also translates to question-answering tasks within high-stakes exams. While LLM hallucinations pose a challenge, their severity and frequency may be less problematic than initially presumed - at least in exam contexts. LLM hallucination is more of an impediment in exams where questions require a faithful recall of facts from subject areas that are underrepresented in the LLM training datasets [40], which is not an issue in most cases. Thus, the impact of hallucination on cheating depends on the nature of the assessment and the narrowness of a subject domain. Highly specialized exams on more esoteric topics which prioritize recall and memorization may be more vulnerable, as plausible-sounding fabrications might go undetected by students. Conversely, tasks requiring in-depth syntheses, deeper reasoning, or source verification present greater challenges for LLMs [26,43]. It is worthwhile acknowledging that cheating often aims to satisfy minimal requirements rather than achieve absolute correctness. This means cheaters often seek to produce work that appears sufficiently knowledgeable to pass, rather than striving for in-depth understanding or absolute accuracy. This perspective is essential when considering AI-assisted cheating. LLMs excel at generating plausible and sometimes merely superficially correct output, aligning with the needs of those seeking to bypass genuine learning. Poorly designed assessments emphasising rote memorisation, simplistic short answers or basic procedural knowledge are particularly susceptible as LLMs can now sufficiently fulfil these minimal requirements. Even if the possibility of LLM hallucination is deemed to be likely in certain assessments, recent studies demonstrate that with careful strategies and prompting techniques, hallucination can to a large degree be mitigated, though not entirely eliminated [18,44–47].

2.3. Reevaluating the Critiques of the Reasoning Capabilities of LLMs

Perceptions that advanced LLMs exhibit limited reasoning abilities still prevail [25,26,43]. This is also maintained across numerous studies exploring LLMs' complex reasoning capabilities in handling academic content. Yeadon and Halliday [48] suggest that while LLMs can function adequately on elementary physics questions, they falter with more advanced content, novel methodologies not included in standard curricula, and basic computational errors. Singla [49] noted that while LLMs like GPT-4 demonstrate proficiency in text-based Python programming, they struggle significantly with visual programming tasks that require a synthesis of spatial, logical, and programming skills. Frequently, LLMs have been observed to under-perform in tasks that necessitate a deep integration of diverse cognitive skills, especially in math-intensive subjects across various languages [36].

However, while there are valid concerns regarding LLMs' current limitations in handling complex reasoning tasks, there is also accumulating evidence of their improving capabilities [32,50,51]. Within an academic context, Liévin *et al.* [52] conclude that LLMs can effectively answer and reason about medical questions, while a recent survey Chang *et al.* [24] indicates that LLMs perform well in tasks like arithmetic reasoning and demonstrate marked competence in logical reasoning tasks too; though, they do encounter significant challenges with abstract and multi-hop reasoning, struggling particularly with tasks requiring complex, novel, or counterfactual thinking. The ability to self-critique is necessary for advanced reasoning that supports rational decision-making and problem-solving, and [53] demonstrate the difficulty of achieving this within LLMs; however, they show how an improvement in LLM's performances on reasoning tasks can be elicited through advanced prompting techniques involving self-critique.

Indeed, studies are beginning to demonstrate how LLMs are becoming more competent at reasoning, and they are also illustrating how effective reasoning can be elicited from LLMs through more sophisticated strategies and prompting techniques that are particularly useful for complex tasks requiring multi-step approaches that decompose the problems into sub-tasks. The next frontier of reasoning complexity is the ability of LLMs to reason across multiple modalities of inputs, where the inputs comprise both text and visualations, and will eventually include additional multimedia inputs. Research in this space has only just begun to emerge and has consistently been indicating limitations of LLMs in their ability to integrate multimodal reasoning accurately. While multi-step reasoning involves a sequential process to reach a conclusion, building logically on each step, multi-hop reasoning entails making several inferential leaps among unlinked data points or different modalities, to piece together an answer. Feng *et al.* [54] concluded that GPT-4 struggled to retain and process visual information in combination with textual inputs that requires multi-hop reasoning, and likewise, [41] found that LLMs generally had variable performances in visual and multimodal medical question-answering tasks, with some under-performing significantly and exhibiting specific deficits, particularly in areas requiring intricate reasoning in medical imaging. Similar findings were arrived at by Stribling *et al.* [32], who assessed the capability of GPT-4 to answer questions from nine graduate-level final examinations in the biomedical sciences where GPT-4 again performed poorly on questions based on figures. Given the research interest in this area [36], created the first benchmark dataset for evaluating LLMs on multimodal multiple-choice questions, concluding that state-of-the-art LLMs struggled with interpreting complex image details in university-level exam questions and overall performed poorly. However, all studies exploring LLMs' multimodal capabilities did not consider multi-hop decomposition strategies that attempt to access the latent sophisticated reasoning capabilities of the LLMs.

2.4. Summary of Literature and Identification of Research Gaps

Current research indicates that concerns regarding LLMs' tendency to produce hallucinated content, while recognised as a limitation, cannot significantly deter their use for effective cheating in most instances. Moreover, while LLMs exhibit certain limitations in complex reasoning tasks, there is clear emerging evidence of their enhanced capabilities, particularly when effectively prompted through

more sophisticated strategies. These strategies, which involve decomposing complex textual tasks into simpler, multi-step stages, have proven effective in eliciting higher-order reasoning from LLMs. LLMs excel in tasks that leverage their strengths in pattern recognition and structured problem-solving but have recently been shown to face significant challenges with reasoning tasks across several modalities of inputs. These currently involve integrating visual and textual inputs in which current models show notable weaknesses in contexts requiring intricate multi-hop reasoning in complex academic testing. A clear gap therefore exists in the development of approaches seeking to improve LLM reasoning on multimodal tasks that demand multi-hop problem-solving strategies.

3. Multimodal LLM Self-Reflective Strategy

This study proposes a multimodal self-reflective strategy for LLMs to demonstrate how to invoke critical thinking and higher-order reasoning within the multimodal LLM agents to steer them towards accurately responding to complex exam questions that integrate textual and visual information, but which can, in theory, be expanded to other modalities as well. The strategy follows a structured iterative process:

1. **Initial Response Evaluation:** The LLM initially responds to a multimodal question, providing a baseline for its capability in interpreting, analysing and integrating all the provided modalities.
2. **Conceptual Self-Reflection:** The LLM is then prompted to assess its own understanding of the key concepts within the textual content of the exam question, inviting the LLM to describe and explain in greater detail its knowledge and understanding of a key concept in isolation from information contained in other modalities.
3. **Visual Self-Reflection:** The LLM is prompted to focus its perception and understanding on key visual cues and elements within an image and to reflect on this, again in isolation from information contained in other modalities.
4. **Synthesis of Self-reflection:** The LLM model is then prompted to reevaluate and revise its initial response from the first step in light of its responses arising from the self-reflective prompting.
5. **Final Response Generation:** The LLM responds with a final answer that uses high-order reasoning to integrate and potentially correct the initial response with a revised answer, aiming for greater accuracy, completeness and depth of understanding.

The proposed iterative self-reflexive strategy builds upon prior works. It can be seen as an expansion of similar interactive self-reflection approaches designed to mitigate hallucination in text-only inputs like that of [18], and self-familiarity of Luo *et al.* [45], as well as the self-critique approach by Luo *et al.* [53] developed to enhance critical thinking in LLMs; alongside chain-of-thought (CoT) techniques [47] that attempt to elicit LLMs to explain their reasoning step-by-step. The proposed multimodal strategy differs from prior techniques, first and foremost by being a multimodal approach, and secondly, it is more prescriptive in how it directs self-reflection.

This approach prompts an LLM in a step-wise manner to revisit and refine its initial responses through a deliberate, guided internal dialogue across conceptual and visual dimensions that constitutes a metacognitive process. This iterative approach that emphasises critical analysis and the synthesis of information aligns with educational principles of self-regulated learning and reflective practice, and the ability of LLM to expand upon, integrate, and critically assess various inputs that mirror higher-order cognitive processes such as those outlined in Bloom's Taxonomy [55].

4. Materials and Methods

The methodology adopted to evaluate the reasoning abilities of GPT-4 across multiple modalities involves two parts, with GPT-4 being selected due to its presently superior multimodal capabilities compared to alternative models [36]. The first part is a case study that tests the proposed multimodal self-reflection strategy on actual university-level exam questions. The second part involved a quantitative and qualitative self-evaluation by GPT-4 on its estimated ability to answer exam questions that comprise textual descriptions of visualisations.

4.1. Evaluation of the multimodal self-reflection strategy

The initial phase of the methodology involved conducting detailed case studies to test the proposed multimodal self-reflection strategy. A university-level exam question from the field of Finance and another from Computer Science to served as case studies. Each question was designed to invoke a high level of reasoning across both textual and visual representations from students, thus allowing for an observation of how GPT-4 applies its multimodal reasoning capabilities in real-world scenarios, and how the proposed strategy can guide the LLM towards a correct answer through self-reflection. Subject experts were used to evaluate the correctness of the responses[1].

4.2. Comprehensive Multimodal Question Assessment

Beyond the case studies, the methodology extended to a broader evaluation involving a dataset of 600 multimodal exam questions. These questions were crafted with the assistance of GPT-4, to cover 12 academic subjects, with the dataset of the questions and responses made available publicly[2]. In this set of experiments, the exam questions were in a text-only format, describing the nature of each exam question together with a description of figures, graphs, diagrams, images, charts, and tables. This approach followed other similar approaches from literature [49,54], whereby instead of directly providing images embedded in the questions themselves, they instead provided descriptions of those images alongside the question itself to GPT-4 for a response.

Each question was designed to mirror the complexity and scope encountered in university-level academic examinations, thus providing an estimated measure of GPT-4's visual and text-based reasoning proficiency. The questions spanned diverse topics within each selected subject, such as Computer Science, Engineering, Nursing, Biology, History, Communications, Education, Marketing, Finance, Economics, and Business Administration/Management, covering the disciplines of Business, Sciences and Humanities with an equal distribution of subjects.

4.3. Proficiency Evaluation Process

GPT-4's proficiency in handling these multimodal questions was evaluated through a dual-phase process. Initially, GPT-4 performed a self-assessment of its ability to respond accurately to each question, rating its proficiency on a scale from 0 to 100. This self-assessment phase allowed the LLM's self-perceived understanding and its ability to analyze and respond to complex multimodal data to be gauged. For each response, GPT-4 was asked to re-evaluate its initial response to arrive at the final proficiency scores. GPT-4 was also tasked with critically analyzing its capability to answer each description of a multimodal question and to explain what the model's current training capacity finds challenging and easy to answer within each question. A selection of questions, proficiency scores and a self-assessment analysis by GPT-4 is shown in Table 1.

---

[1] The author of this paper was one of the evaluators since the Computer Science question was devised by the author and was originally used in an actual exam to evaluate students.
[2] https://github.com/teosusnjak/multimodal-chatgpt-exam-evaluation

Table 1. Examples of three descriptive multimodal exam questions, including GPT-4's self-assessment of its ability to answer the question and quantitative proficiency scores evaluating its ability to respond accurately.

Question Type (Subject)	Question and Competence Self-evaluation	Proficiency Score
(Nursing) Surgical Outcomes Dashboard	**Question:** Analyze a dashboard displaying surgical outcomes, including success rates, complication rates, and patient satisfaction scores. Discuss how this data informs surgical quality improvement. **Self-Analysis:** GPT-4 could struggle with evaluating surgical outcomes data, particularly in understanding how success rates, complication rates, and patient satisfaction scores inform surgical quality improvement efforts.	70
(Education) Educational Technology Tools Comparison Table	**Question:** Evaluate a table comparing various educational technology tools based on functionality, usability, and cost. Discuss how these tools can be effectively integrated into the classroom. **Self-Analysis:** While GPT-4 can compare educational technology tools, fully grasping how these tools integrate into classroom settings and enhance learning requires knowledge of pedagogical practices and technology effectiveness in education.	83
(Engineering) Automotive Aerodynamics Simulation Results	**Question:** Evaluate the results of computational fluid dynamics simulations for car models. Discuss how aerodynamics influences vehicle design and fuel efficiency. **Self-Analysis:** GPT-4 can evaluate aerodynamics simulation results but may have difficulty in fully understanding the implications for vehicle design and fuel efficiency, which require integration of fluid dynamics principles and automotive engineering practices.	80

5. Results

The presentation of results begins with the illustration of the proposed method to invoke iterative self-reflection within the multimodal LLMs, with the goal of triggering and demonstrating advanced reasoning capabilities inherent within the selected LLM. Two illustrations of the multimodal self-reflection strategy are shown on exam questions from Finance and Computer Science respectively. The second part of the results section covers the results of GPT-4's estimation of its capabilities to answer multimodal exam questions across the 12 subject areas.

5.1. Case Study - Finance

In the first example, a multiple-choice exam question from a Finance course is taken as seen in Figure 1. Initially, a baseline was established to see if GPT-4 with vision could answer the question correctly when offered multiple choices of responses to choose from. The correct answer is "B", and GPT-4 correctly answered this, with the response: "The most accurate interpretation based on the shape of the yield curve would be option B, indicating that the market expects inflation to fall sharply in the short-term.". With this established, the next aim was to demonstrate the efficacy of the proposed approach. Table 2 shows the Steps (1 to 4) that were carried out in a new and separate session. GPT-4 was first asked to answer the exam question without providing multiple-choice answers to test its ability to fully combine visual and conceptual reasoning, to which the response given was incorrect (Step 1). Subsequently, GPT-4 was asked to self-reflect and describe the relevant concepts (Step 2) and to self-reflect on its perception of important visual cues (Step 3) in the image. Both exercises produced sufficiently correct reflective responses, upon which, the exam question was again posed (Step 4) with GPT-4 being invited to revise its initial response based on its subsequent reflective responses. The culmination of these exercises produced the correct final answer.

Table 2. Finance exam question case study example illustrating the step-wise invocation of self-reflection within GPT-4 that is concept and visual-reasoning based. Key aspects of the LLM responses are highlighted.

Step	User Prompt	GPT-4 Vision Response
(1) Initial response requiring concept and image reasoning	"Answer the following Finance exam question based on the supplied image. What trend will the market expect that the inflation will follow?"	"The market expectation inferred from the yield curve in the image is that in the short term, inflation may rise or remain relatively high, prompting central banks to maintain or raise interest rates to control inflationary pressures. Conversely, in the long term, the expectation of lower yields suggests that the market anticipates a reduction in inflation, potentially due to economic slowdown or effective monetary policy measures leading to lower interest rates." INCORRECT ANSWER
(2) Self-reflection focusing on concept reasoning	"Expand and describe the expectations theory in the context of yield curves and how they affect short-term and long-term inflation."	"The expectations theory holds that the yield curve affects inflation expectations in both the short and long term, although the mechanisms and strengths of these effects may differ. Short-term rates are closely tied to the central bank's policy rate, which is influenced by the central bank's inflation target. If the market believes that inflation will rise above the central bank's target, it expects the central bank to increase the policy rate to combat inflation, thus increasing short-term yields. Consequently, a steep yield curve can signal rising short-term inflation expectations. The long end of the curve reflects the market's expectations for inflation over a longer horizon. Suppose long-term rates are low relative to short-term rates, creating an inverted curve. In that case, it can indicate that the market expects the central bank to eventually reduce policy rates in response to lower inflation or even deflationary pressures. This scenario would indicate a forecast of stable to declining long-term inflation."[3]CORRECT
(3) Self-reflection focusing on the image reasoning	"Based on the yield curve in the image, at what point in time is the change in yield most acute?"	"The change in yield is most acute at the shorter end of the spectrum, particularly between the 3-month and 2-year maturities. This portion of the curve displays a steeper slope compared to the more gradual slope between the 2-year and 10-year maturities." CORRECT
(4) Revised response post-self-reflection	"Based on your observations in the response and reasoning above, revise and improve your answer to the above question about the effect of yield curves on short-term and long-term inflation."	"Based on the steepness of the yield curve between the 3-month and 2-year maturities, the market appears to expect a relatively sharp decline in short-term inflation expectations. The leveling off of the yield curve beyond the 2-year mark suggests that the market anticipates more stable, modest changes to long-term inflation expectations." CORRECT ANSWER

---

3   The statement while largely correct, does oversimplify the relationship between the yield curve and inflation expectations by focusing narrowly on inflation and central bank policies, while ignoring other influential factors and the complexity of economic decision-making. Nonetheless, the oversimplification is sufficient overall to generate the correct final response.
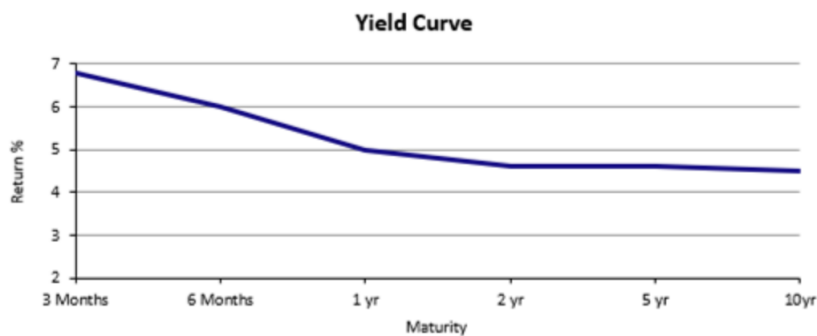
Yield Curve

Q7.   Based on the expectations theory you could deduce from the **<u>above</u>** yield curve that the market believes inflation will:
- A)   Remain at current levels over the next ten years
- B)   Fall sharply in the short-term *
- C)   Initially fall in the short-term and then rise again in the long-tem
- D)   Fall sharply in the long-term
- E)   None of the above

Figure 1. An example of a multiple-choice question from a Finance exam.

The metacognitive process illustrated in Table 2, in the form of self-reflection is being invoked within an LLM and evidenced by the model's ability to engage with questions that challenge its understanding, with an openness to revising its perspective based on new insights or information. A deliberate and structured review of the model's initial responses was demonstrated taking place at conceptual and visual perception levels. This iterative process with the emphasis on revision and improvement of understanding, based on a guided internal dialogue, captured the essence of self-reflection. Self-reflection is also intrinsically linked to critical thinking and higher-order reasoning which underscores the advanced and latent capabilities of GPT-4. Critical thinking involves the objective analysis and evaluation of an issue in order to form a judgment, meanwhile, higher-order reasoning is characterised by the ability to understand complex concepts, apply multiple concepts simultaneously, analyze information, synthesize insights from various sources, and evaluate outcomes and approaches. In the example, the self-reflective process undertaken by the model to "expand and describe" a relevant concept, engages in a form of analysis which is one of the hallmarks of critical thinking. Meanwhile, GPT-4's higher-order reasoning was demonstrated when the model was prompted to integrate and synthesise separate reflections on concepts and visual information to form a comprehensive understanding of the phenomenon in question. This represents the synthesis and evaluation aspects of Bloom's Taxonomy, which are considered higher-order cognitive skills.

5.2. Case Study - Computer Science

The same analysis was replicated on a university-level exam question from Computer Science. This was a short answer question requiring critical reasoning, interpretation of visual patterns and their reconciliation with theoretical concepts. The question can be seen in Figure and the initial incorrect response, as well as the multimodal invocation of the self-reflective strategy, can be seen in Table 3.

Table 3. Machine learning exam question case study example illustrating the step-wise invocation of self-reflection within GPT-4 that is concept and visual-reasoning based. Key aspects of the LLM responses are highlighted.

Step	User Prompt	GPT-4 Vision Response
(1) Initial response requiring concept and image reasoning	"Examine the figure below, which illustrates the convergence behaviour of training and generalization error for a hypothetical machine learning classifier trained by AdaBoost under typical conditions. Discuss the plausibility of observing these kinds of patterns with an algorithm like AdaBoost."	"The figure presented demonstrates a typical convergence pattern for AdaBoost, a popular ensemble learning algorithm...The convergence behavior depicted in the graph is thus a plausible outcome when using AdaBoost without preventive measures against overfitting." INCORRECT ANSWER
(2) Self-reflection focusing on concept reasoning	"Tell me about the typical effects on generalization by an algorithm like AdaBoost and, under normal circumstances how much it is likely to compromise generalizability by overfitting. "	"Under standard conditions, the algorithm's strength lies in its adaptive nature, concentrating on instances that were previously misclassified and thus, when properly tuned, it can be less prone to overfitting compared to non-ensemble methods. However, overfitting with AdaBoost can occur in the presence of noise and outliers in the training data...In a controlled experimental setting, it is often possible to achieve generalization error rates that are competitive with or even superior to other algorithms, provided that AdaBoost's parameters are optimized through techniques like cross-validation. However, in the presence of highly noisy data or when the assumptions underlying the algorithm are violated, the deterioration in generalization could be considerable. Therefore, while AdaBoost has the potential for robust generalization, it also has a non-negligible risk of compromising generalizability due to overfitting, especially when the model complexity is not adequately controlled." CORRECT
(3) Self-reflection focusing on the image reasoning	"Examine the figure carefully noting the value of the error when the overfitting you mentioned begins to take place and note carefully the point it has reached."	"The graph indicates that overfitting begins to occur just after 10 boosting rounds, where the generalization error reaches its nadir and starts to increase. At this juncture, the training error continues to decrease, demonstrating the model's increasing fit to the training data. By the 40th boosting round, the generalization error has increased significantly, surpassing an error rate of 0.4, which is a marked rise from its lowest point. This reflects a substantial degradation in the model's predictive accuracy on unseen data, signaling a pronounced overfitting as the model is overly optimized for the training dataset." CORRECT
(4) Revised response post-self-reflection	"Now, consider your response above about AdaBoost's effects on generalisability and combine your response with the detailed observations above regarding when overfitting started and the significant degree to which it increased. Based on these responses, revise your answer to the original exam question about the plausibility of observing these trends with AdaBoost under standard conditions where the data is not corrupted and everything else is typical."	"Upon integrating the detailed observations with the known effects of AdaBoost on generalizability, it is pertinent to re-evaluate the plausibility of the presented trends under typical, controlled conditions. When the dataset is clean and representative of the underlying distribution, AdaBoost, by incrementally correcting errors of the weak learners, should ideally show a decrease in both training and generalization errors. The convergence pattern would typically demonstrate a plateauing of the generalization error as the number of boosting rounds increases, indicating that the model has learned the general trends in the data without being unduly influenced by noise... Therefore, in a standard scenario with typical data quality and distribution, the observed pattern of the generalization error increasing sharply after a certain point would not be considered typical for AdaBoost...In conclusion, the graph's indication of rapid overfitting after a relatively low number of boosting rounds is implausible in standard conditions with AdaBoost, where data is clean and well-prepared." CORRECT ANSWER

Examine the figure below, which illustrates the convergence behaviour of training and generalization error for a hypothetical machine learning classifier trained by AdaBoost under typical conditions. Discuss the plausibility of observing these kinds of patterns with an algorithm like AdaBoost.
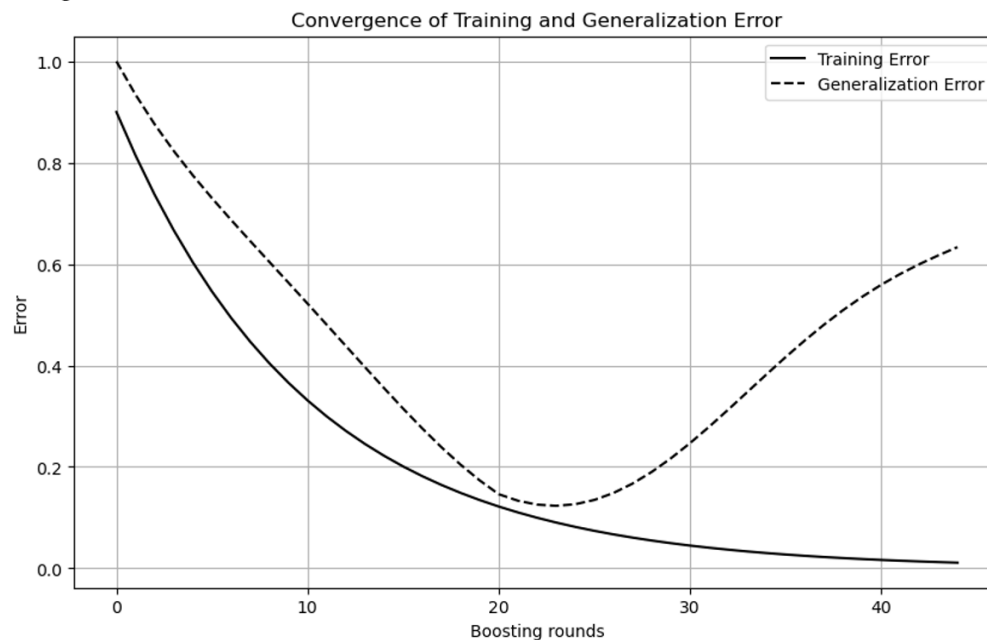

Figure 2. An example of a multimodal Computer Science exam question.

In line with the first case study example, the same pattern can be observed. The task of fully understanding the question, then integrating critical reasoning with the perceptive reasoning of a visual artefact, alongside the reconciliation of the observations with theoretical concepts, is too complex for multimodal LLMs to perform simultaneously (Step 1); confirming findings from other studies [24,53]. However, when the complexity is reduced and isolated to self-reflective exercises on individual modalities in turn (Steps 2 and 3), and then combined, the initial incorrect response is revised with a correct answer.

5.3. GPT-4 Multimodal Capability Estimations by Subject Area

Here, a broader and more general estimation of GPT-4's multimodal capabilities for answering exam questions is presented. More specifically, GPT-4's own estimations of its proficiency to answer 600 multimodal exam questions across 12 subject areas are quantified. The analysis of the means and standard deviations of the estimated proficiency scores can be seen in Figures 3 and 4 respectively indicating GPT-4's assessment of its competence to contextualise, visually perceive, reason about, and generate coherent responses based on multimodal inputs. When evaluating the mean proficiency scores, the disciplines of biology, history, and psychology rank highest, suggesting GPT-4 exhibits a stronger alignment with the types of visual information and analytical reasoning these fields typically employ. This could be due to the rich contextual cues present in visual materials like biological diagrams or historical timelines, which offer structured and often hierarchical information that aligns well with GPT-4's training on pattern recognition and sequence alignment. In contrast, disciplines such as business administration and economics, which often involve complex and abstract data, may require a deeper understanding of human and market behaviours which currently present challenges, and are evidences in lower mean proficiency scores.
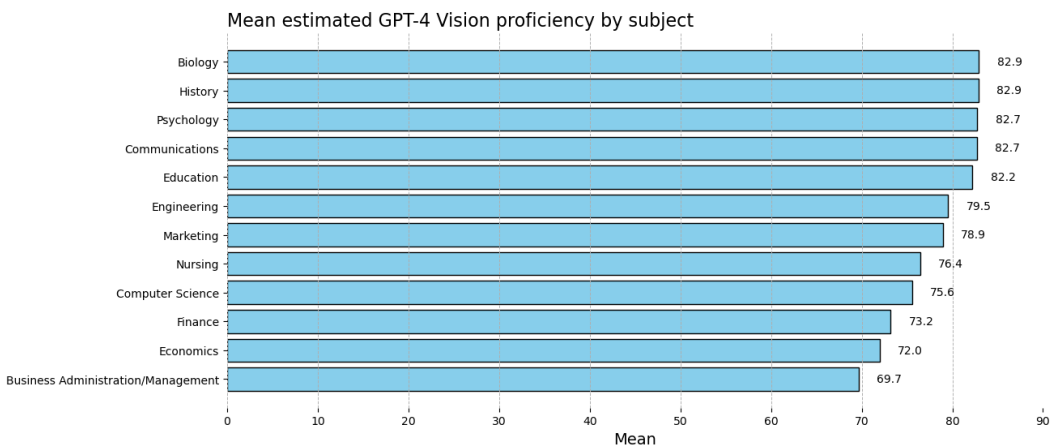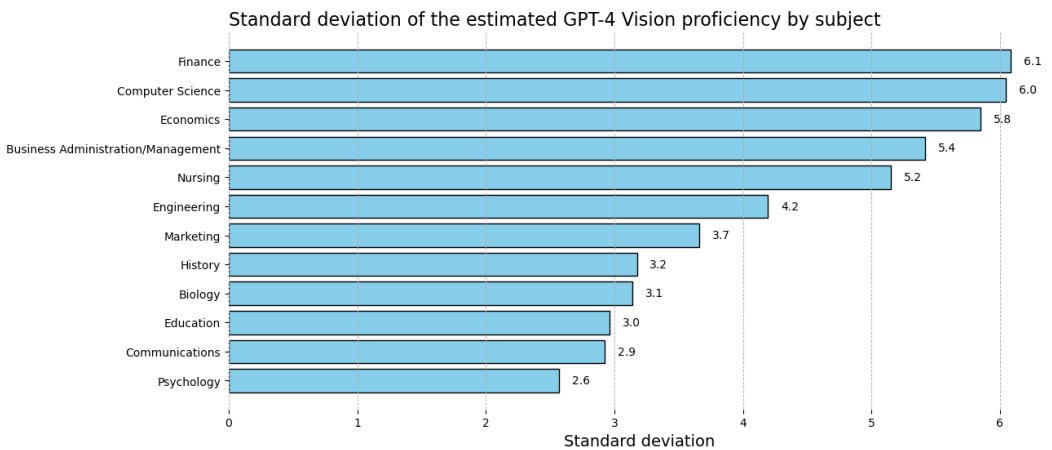
Figure 3. Mean.

Figure 4. Standard deviation.

The standard deviation scores indicate the consistency of GPT-4's performance across subjects. Higher variability in fields like finance and computer science suggests that GPT-4's understanding may fluctuate significantly depending on the specificity of the task or the complexity of the visual data. This could imply that while GPT-4 can proficiently handle standard multimodal questions within these domains, it might struggle with more complex or less conventional topics, or those requiring deeper inferential reasoning. Moreover, the relatively lower variances observed in psychology and communications imply a more uniform proficiency across different queries within these subjects. This may be attributed to the nature of data in these fields, which frequently include human behavioural patterns and communicative structures—areas where GPT-4 has substantial experience from training datasets. Table 4 collates all the analyses across 600 responses and identifies features of multimodal exam questions per subject that GPT-4 perceives via self-assessment, to be able to handle both with high proficiency as well as with some degree of difficulty.

Table 4. An overarching summary of all the exam question performance evaluations of GPT4's strengths and limitations across all subjects.

Subject	GPT-4 Vision's performance characteristics
Business Administration/ Management	Excels in extracting and summarizing business-related information from visual data, recognizing patterns and trends. Exhibits challenges in grasping the full scope of socio-economic contexts and generating deeper strategic insights.
Finance	Displays adeptness at quantitative financial analysis and understanding fundamental financial concepts. It encounters challenges with complex financial theories and strategic real-world applications.
Marketing	Excels in interpreting marketing data and grasping core concepts, aiding in trend identification. It struggles with deciphering more complex strategic implications and nuances of consumer psychology.
Economics	Identifies trends from economic data and understands foundational principles, but faces limitations in deeper theoretical analyses and predictive economic impacts.
Computer Science	Interprets technical diagrams and data trends in computer science contexts effectively but struggles with more sophisticated system dynamics and predictive analysis.
Nursing	Adeptly interprets data and fundamental nursing concepts but faces challenges in more demanding clinical reasoning and holistic healthcare strategy development.
Engineering	Excels in parsing engineering data and explaining technical concepts but encounters difficulties with contextual analyses and predictive evaluations.
Biology	Shows proficiency in interpreting biological data and explaining processes but struggles with understanding more complex concepts and performing predictive analysis.
Education	Parses educational data and links theories to practice well but struggles with the complexities of educational systems and multidisciplinary integration.
Psychology	Effectively interprets psychological data but struggles with more demanding constructs and forward-looking analyses that require a deeper understanding.
History	Processes historical data and concepts well but struggles with analyzing more complex relationships and conducting critical evaluations.
Communications	Analyzes communication trends and strategies effectively but lacks depth in grasping socio-cultural impacts and strategic ethical considerations.

Finally, all the proficiency scores are aggregated by the overarching disciplines and depicted in Figure 5. In the humanities, which registers the highest proficiency score, the nature of the discipline itself is likely playing a pivotal role. Humanities research and study often encompass a broad spectrum of data interpretation, which points to a close alignment with GPT-4's strengths in language understanding and integration of contextual information. The humanities' interpretive nature allows for a wider margin of acceptable responses, providing GPT-4 with a conducive environment for showcasing its capacity to draw connections between disparate historical events, socio-cultural dynamics, and

philosophical concepts. Furthermore, humanities subjects often demand a high level of narrative construction, which is also well-suited to GPT-4's design that inherently focuses on language and narrative generation.
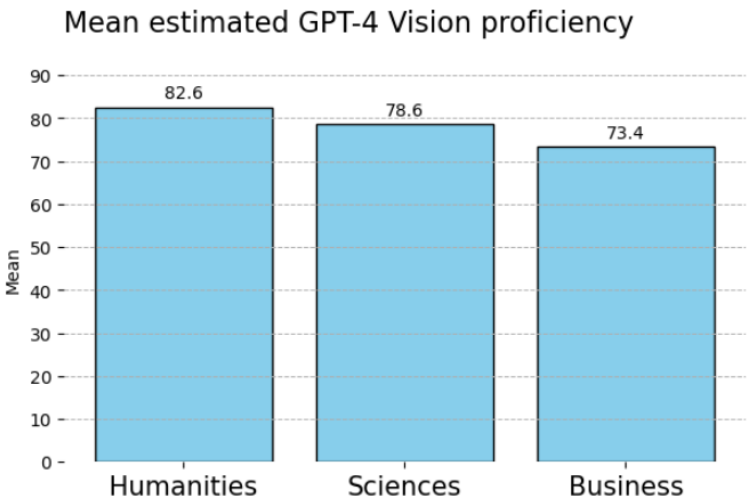

Figure 5. Mean proficiencies by disciplines.

In the sciences, the observed proficiency dip may be attributed to the quantitative and empirical rigidity of the field. Sciences often require precise and unequivocal interpretations, leaving less room for the breadth of interpretative responses that GPT-4 can generate well. Moreover, scientific data frequently necessitate a deeper understanding of causality, experimental design, and statistical validity, which can be challenging for GPT-4 currently, as its abilities are modelled on pattern recognition rather than first principles reasoning. While GPT-4 is competent at identifying patterns and trends in scientific data, the intricacies of scientific theory and the need for detailed methodological analysis can present challenges, which might account for the slightly lower proficiency scores compared to humanities.

For the business discipline, which shows the lowest mean proficiency score, the complexity likely arises from the nature of business decision-making, which often requires an integration of both quantitative data and human judgment. Business subjects not only involve financial and operational data interpretation but also necessitate an understanding of market dynamics, consumer behaviour, and strategic decision-making under uncertainty. These areas rely heavily on real-time data, contextual subtleties, and forward-looking predictions, which are challenging for GPT-4 given its limitations in temporal awareness and predictive modelling based on past and present data trends alone.

6. Discussion

This study has shown the capability of LLMs to answer multimodal exam questions that require advanced reasoning, something that recent studies have identified as a limitation. This work has specifically developed a technique and demonstrated how it can be used to invoke self-reflection within LLMs at different input modalities comprising exam questions, to eventually arrive at correct responses, and thereby, without the intention to do so, this study has provided a "how-to" recipe for more effective cheating on sophisticated exam questions that are unproctored. However, as educators, we need to be aware of these capabilities. A silver lining does exist though since successfully executing a sequence of appropriate and contextualised self-reflective questions that are relevant for each modality in a multimodal exam question, does require some sophistication and level of knowledge about the exam question subject from a person attempting to cheat at least. Additionally, this technique can be used by academics in order to assist with research involving interpretation and reasoning about figures and theoretical concepts.

Nonetheless, it is clear from the results that LLMs possess a noteworthy level of reasoning capability that this study has shown also extends to both text-based and visual modalities. LLMs'

reasoning is improving and will likely continue to do so; however, limits and plateaus to this capability are also reasonable to expect, thus there is a need to continuously probe this capability *vis-a-vis* their proficiency at answering complex multimodal exam questions to determine their upper-performance limits under the current transformer-based [56] machine learning architectures. The iterative self-reflective procedure outlined in this study works, and it has strong theoretical underpinnings that can explain why it works. Transformer-based models that underlie LLMs like GPT-4, comprise two components: the encoder and the decoder. The encoder excels at analyzing and understanding prompts texts while the decoder focuses on producing text based on this understanding. However, the encoder processes the input prompts holistically at once enabling it to capture complex interdependencies, while the decoder is constrained by being able to focus only on previously generated portions of the response without recourse to to consider its future response generation. At the risk of oversimplifying the dynamic, this means that the LLM's ability to understand prompts has certain strengths over its ability to generate responses. Therefore, the encoder's strengths can be leveraged and can be loosely conceptualized as a form of AI self-reflection, where iterative prompting is conducted to extract improved understanding and better responses which are then collectively fed back to the encoder for a final response. The demonstrated process allows the encoder to refine the understanding and context, which the decoder then uses to produce more accurate and contextually appropriate responses. Such iterative, multi-hop decomposition of a complex reasoning task, harnesses the encoder's robust analytical abilities, by progressively improving the quality of the generated responses - and the whole procedure can be viewed as iterative multimodal self-reflection.

This finding has serious implications for online exams. One of the last remaining strategies for making exam questions 'LLM proof' has been to incorporate visual components alongside text. Numerous recent studies have confirmed that this was the case; however, this work has shown that advanced reasoning capabilities are latent within LLMs that can be accessed for solving complex multimodal tasks. It is now incumbent on researchers to apply these types of multi-hop self-reflective techniques on larger multimodal benchmark datasets to more comprehensively quantify the capacities of LLMs on these types of tasks.

With respect to results from estimated multimodal capabilities of GPT-4 across different disciplines, we can infer that GPT-4 shows a particular strength in handling tasks that involve interpretive and descriptive analysis typical of the humanities likely due to having been trained on extensive collections of narrative and descriptive texts which translates to more accurate explanations for visuals related to these subjects. On the other hand, in the sciences and business where the questions may more frequently require predictive, forward-looking and strategic thinking, GPT-4's limitations become more suggested. Its perceived inability to fully grasp real-world time progression and conduct original research and hypothesise about possible outcomes may result in less accurate performance. This observation indicates that GPT-4's effectiveness varies significantly with the nature of the multimodal exam topic, heavily dependent on the type of reasoning the subject demands and the characteristics of the underlying data used in its training. Thus, while GPT-4 estimates broad capabilities of processing a diverse range of visuals, its performance distinctly mirrors the intrinsic features of each academic discipline and the current limitations of AI technology which performs well in identifying patterns and relationships but falls short in understanding causality and making strategic predictions.

Recommendations

Based on the research findings, this work proposes the following strategies that may help in the short term enhance the integrity and effectiveness of online assessments in the context of the advanced reasoning capabilities of LLMs:

1. **Proctored online exams:** There is no substitute to effective proctoring. Therefore, it is recommended to ensure that all online exams are proctored as extensively as possible. Proctoring technologies that include real-time monitoring have their limitations, but they can deter the mis-

use of LLMs and other digital aids. Unproctored exams in the context of existing, and improving multimodal LLM capabilities can no longer be regarded as possessing validity.

2. **Reinstatement of viva-voce exams:** The reintroduction of viva-voce examinations, conducted online, can complement a suite of other assessments. Although viva-voce exams also possess limitations (as do all assessment types), they offer a dynamic, generally reliable, and direct assessment method for measuring student knowledge and reasoning skills. These exams are akin to professional interviews commonly used in industry, thus making them relevant for preparing students for real-world contexts.

3. **Enhanced multimodal exam strategies:** If proctoring of online exams is not feasible, multimodal exams should be designed to maximally increase the cognitive load and processing complexity, making them more challenging for LLMs to provide reliable cheating assistance:

    - Include multiple images alongside text per question to invoke a higher degree of reasoning across the modalities. This approach will increase the complexity of the questions and require a deeper level of reasoning and synthesis, which current LLMs may struggle to manage effectively.
    - Design questions that necessitate the formulation of long-term strategies, forecasts, and projections. These types of questions require not only higher levels of conceptual understanding as well as causal relationships, but also the ability to project future trends and consequences which this research shows are a challenge to the predictive capabilities of LLMs.
    - Integrate real-world scenarios that are current and relevant. Questions that reflect very recent developments or ongoing complex real-world problems and require up-to-date knowledge, making it difficult for LLMs to reason and generate accurate responses based solely on pre-existing and limited data for certain topics.
    - Consider incorporating additional modalities into the questions such as video-based and/or audio-based questions alongside images and text, thus fully exploiting the current limitations of LLMs to process ot incorporate them all simultaneously.
    - Explore formulating multimodal questions that require students to annotate the provided figure(s) or draw an additional figure as part of their answer which would again exploit some of the current LLM limitations.
    - As much as possible, consider ways of linking multimodal exam questions with prior assessments completed by students during a teaching semester, and other course materials which would increase the difficulty of the LLM in producing a correct response.
    - Consider incorporating some decoy questions specifically designed to detect LLM assistance. These questions could be subtly designed to prompt LLMs into revealing their non-human reasoning patterns through specific traps that exploit known LLM weaknesses, such as generating responses based on unlikely combinations of concepts or unusual context switches that a human would likely not make. This would not however be straightforward to implement since different multimodal LLMs will likely also have different responses to the decoy questions.

Limitations and Future Work

This study makes strides in utilising LLMs for enhancing academic assessment but recognises the need for broader empirical support, suggesting areas for future exploration. This work was limited by only two multimodal exam case studies and further research should extend this approach across more disciplines and question types, particularly focusing on developing a benchmark dataset comprising deep reasoning multimodal exam questions and not merely multiple choice questions that currently do exist. Moreover, while the evaluations of exam questions by GPT-4 provide a useful gauge of the model's current capabilities, they remain estimates and necessitate more rigorous, quantitative validation to accurately measure its actual performance. Additionally, the study highlights the manual nature of the self-reflective process used to trigger latent LLM reasoning, and an algorithmic

automation of the proposed strategy that can be applied across multiple modalities to improve the scalability of the approach should be explored and developed.

7. Conclusions

This study critically addresses the role of Large Language Models (LLMs) like ChatGPT in modern online educational exams, highlighting the challenges they pose to academic integrity in the absence of proctoring. Key contributions of this research include the introduction of a novel iterative strategy that invokes self-reflection within LLMs to guide them toward correct responses to complex multimodal exam questions, thereby demonstrating multi-hop reasoning capabilities latent within LLMs. By invoking self-reflection within LLMs on each separate modality, the proposed strategy demonstrated how critical thinking and higher-order reasoning can be triggered and integrated in a step-wise manner to steer LLMs towards correct answers, demonstrating that LLMs can be used effectively for cheating on multimodal exam questions. The study also conducted a broad evaluation using descriptions of 600 multimodal exam questions across 12 university subjects, to estimate the proficiency of LLMs with vision capabilities to answer university-level exam questions involving visuals, suggesting that exam questions including visuals from humanities may pose the least amount of challenge to answer correctly by best-performing LLMs, followed by a question from the sciences and business respectively. Additionally, this work offers pragmatic recommendations for conducting exams and formulating exam questions in view of the current reasoning and multimodal capabilities of LLMs.

References

1. Barber, M.; Bird, L.; Fleming, J.; Titterington-Giles, E.; Edwards, E.; Leyland, C. Gravity assist: Propelling higher education towards a brighter future - Office for Students, 2021.
2. Butler-Henderson, K.; Crawford, J. A systematic review of online examinations: A pedagogical innovation for scalable authentication and integrity. *Computers & Education* **2020**, *159*, 104024.
3. Coghlan, S.; Miller, T.; Paterson, J. Good proctor or "big brother"? Ethics of online exam supervision technologies. *Philosophy & Technology* **2021**, *34*, 1581–1606.
4. Henderson, M.; Chung, J.; Awdry, R.; Mundy, M.; Bryant, M.; Ashford, C.; Ryan, K. Factors associated with online examination cheating. *Assessment & Evaluation in Higher Education* **2022**, pp. 1–15.
5. Dumulescu, D.; Muţiu, A.I. Academic leadership in the time of COVID-19—Experiences and perspectives. *Frontiers in Psychology* **2021**, *12*, 648344.
6. Whisenhunt, B.L.; Cathey, C.L.; Hudson, D.L.; Needy, L.M. Maximizing learning while minimizing cheating: New evidence and advice for online multiple-choice exams. *Scholarship of Teaching and Learning in Psychology* **2022**.
7. Garg, M.; Goel, A. A systematic literature review on online assessment security: Current challenges and integrity strategies. *Computers & Security* **2022**, *113*, 102544.
8. Arnold, I.J. Cheating at online formative tests: Does it pay off? *The Internet and Higher Education* **2016**, *29*, 98–106.
9. Ahsan, K.; Akbar, S.; Kam, B. Contract cheating in higher education: a systematic literature review and future research agenda. *Assessment & Evaluation in Higher Education* **2021**, pp. 1–17.
10. Crook, C.; Nixon, E. How internet essay mill websites portray the student experience of higher education. *The Internet and Higher Education* **2021**, *48*, 100775.
11. Noorbehbahani, F.; Mohammadi, A.; Aminazadeh, M. A systematic review of research on cheating in online exams from 2010 to 2021. *Education and Information Technologies* **2022**, pp. 1–48.
12. Allen, S.E.; Kizilcec, R.F. A systemic model of academic (mis)conduct to curb cheating in higher education. *Higher Education* **2023**, pp. 1–21.

13.   Henderson, M.; Chung, J.; Awdry, R.; Ashford, C.; Bryant, M.; Mundy, M.; Ryan, K. The temptation to cheat in online exams: moving beyond the binary discourse of cheating and not cheating. *International Journal for Educational Integrity* **2023**, *19*, 21.

14.   Malik, A.A.; Hassan, M.; Rizwan, M.; Mushtaque, I.; Lak, T.A.; Hussain, M. Impact of academic cheating and perceived online learning effectiveness on academic performance during the COVID-19 pandemic among Pakistani students. *Frontiers in Psychology* **2023**, *14*, 1124095.

15.   Newton, P.M.; Essex, K. How common is cheating in online exams and did it increase during the COVID-19 pandemic? A systematic review. *Journal of Academic Ethics* **2023**, pp. 1–21.

16.   McCabe, D.L. CAI Research Center for Academic Integrity, Durham, NC, 2005.

17.   Wajda-Johnston, V.A.; Handal, P.J.; Brawer, P.A.; Fabricatore, A.N. Academic dishonesty at the graduate level. *Ethics & Behavior* **2001**, *11*, 287–305.

18.   Ji, Z.; Yu, T.; Xu, Y.; Lee, N.; Ishii, E.; Fung, P. Towards Mitigating Hallucination in Large Language Models via Self-Reflection, 2023, [arXiv:cs.CL/2310.06271].

19.   Richards, M.; Waugh, K.; Slaymaker, M.; Petre, M.; Woodthorpe, J.; Gooch, D. Bob or Bot: Exploring ChatGPT's Answers to University Computer Science Assessment **2024**. *24*. doi:10.1145/3633287.

20.   Chan, J.C.; Ahn, D. Unproctored online exams provide meaningful assessment of student learning. *Proceedings of the National Academy of Sciences* **2023**, *120*, e2302020120.

21.   Martino, A.; Iannelli, M.; Truong, C. Knowledge injection to counter large language model (LLM) hallucination. European Semantic Web Conference. Springer, 2023, pp. 182–185.

22.   Yao, J.Y.; Ning, K.P.; Liu, Z.H.; Ning, M.N.; Yuan, L. LLM Lies: Hallucinations are not Bugs, but Features as Adversarial Examples, 2023, [arXiv:cs.CL/2310.01469].

23.   Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; Wang, L.; Luu, A.T.; Bi, W.; Shi, F.; Shi, S. Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models, 2023, [arXiv:cs.CL/2309.01219].

24.   Chang, Y.C.; Wang, X.; Wang, J.; Wu, Y.; Zhu, K.; Chen, H.; Yang, L.; Yi, X.; Wang, C.; Wang, Y.; Ye, W.; Zhang, Y.; Chang, Y.; Yu, P.S.; Yang, Q.; Xie, X. A Survey on Evaluation of Large Language Models. *ArXiv* **2023**, *abs/2307.03109*. doi:10.48550/arXiv.2307.03109.

25.   McKenna, N.; Li, T.; Cheng, L.; Hosseini, M.J.; Johnson, M.; Steedman, M. Sources of Hallucination by Large Language Models on Inference Tasks **2023**. pp. 2758–2774. doi:10.48550/arXiv.2305.14552.

26.   Liu, H.; Ning, R.; Teng, Z.; Liu, J.; Zhou, Q.; Zhang, Y. Evaluating the Logical Reasoning Ability of ChatGPT and GPT-4. *ArXiv* **2023**, *abs/2304.03439*. doi:10.48550/arXiv.2304.03439.

27.   Schultz, M.; Callahan, D.L. Perils and promise of online exams. *Nature Reviews Chemistry* **2022**, *6*, 299–300.

28.   Cotton, D.R.; Cotton, P.A.; Shipway, J.R. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International* **2024**, *61*, 228–239.

29.   Alessio, H.M.; Malay, N.; Maurer, K.; Bailer, A.J.; Rubin, B. Examining the effect of proctoring on online test scores. *Online Learning* **2017**, *21*, 146–161.

30.   Han, S.; Nikou, S.; Ayele, W.Y. Digital proctoring in higher education: a systematic literature review. *International Journal of Educational Management* **2023**, *38*, 265–285.

31.   Abd-Alrazaq, A.; AlSaad, R.; Alhuwail, D.; Ahmed, A.; Healy, P.M.; Latifi, S.; Aziz, S.; Damseh, R.; Alrazak, S.A.; Sheikh, J.; others. Large language models in medical education: opportunities, challenges, and future directions. *JMIR Medical Education* **2023**, *9*, e48291.

32.   Stribling, D.; Xia, Y.; Amer, M.K.; Graim, K.S.; Mulligan, C.J.; Renne, R. The model student: GPT-4 performance on graduate biomedical science exams. *Scientific Reports* **2024**, *14*, 5670.

33.   Rudolph, J.; Tan, S.; Tan, S. ChatGPT: Bullshit spewer or the end of traditional assessments in higher education? *Journal of applied learning and teaching* **2023**, *6*, 342–363.

34.   Lo, C.K. What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences* **2023**, *13*, 410.

35.   Nikolic, S.; Daniel, S.; Haque, R.; Belkina, M.; Hassan, G.M.; Grundy, S.; Lyden, S.; Neal, P.; Sandison, C. ChatGPT versus engineering education assessment: a multidisciplinary and multi-institutional benchmarking and analysis of this generative artificial intelligence tool to investigate assessment integrity. *European Journal of Engineering Education* **2023**, *48*, 559–614.

36. Zhang, W.; Aljunied, M.; Gao, C.; Chia, Y.K.; Bing, L. M3exam A multilingual, multimodal, multilevel benchmark for examining large language models. *Advances in Neural Information Processing Systems* **2024**, *36*.

37. Sadasivan, V.S.; Kumar, A.; Balasubramanian, S.; Wang, W.; Feizi, S. Can AI-Generated Text be Reliably Detected? *ArXiv* **2023**, *abs/2303.11156*. doi:10.48550/arXiv.2303.11156.

38. Orenstrakh, M.S.; Karnalim, O.; Suarez, C.A.; Liut, M. Detecting LLM-Generated Text in Computing Education: A Comparative Study for ChatGPT Cases, 2023, [arXiv:cs.CL/2307.07411].

39. Kumarage, T.; Agrawal, G.; Sheth, P.; Moraffah, R.; Chadha, A.; Garland, J.; Liu, H. A Survey of AI-generated Text Forensic Systems: Detection, Attribution, and Characterization, 2024, [arXiv:cs.CL/2403.01152].

40. Kalai, A.T.; Vempala, S.S. Calibrated Language Models Must Hallucinate, 2024, [arXiv:cs.CL/2311.14648].

41. Pal, A.; Sankarasubbu, M. Gemini Goes to Med School: Exploring the Capabilities of Multimodal Large Language Models on Medical Challenge Problems & Hallucinations, 2024, [arXiv:cs.CL/2402.07023].

42. Nori, H.; King, N.; McKinney, S.; Carignan, D.; Horvitz, E. Capabilities of GPT-4 on Medical Challenge Problems. *ArXiv* **2023**, *abs/2303.13375*. doi:10.48550/arXiv.2303.13375.

43. Stechly, K.; Marquez, M.; Kambhampati, S. GPT-4 Doesn't Know It's Wrong: An Analysis of Iterative Prompting for Reasoning Problems. *ArXiv* **2023**, *abs/2310.12397*. doi:10.48550/arXiv.2310.12397.

44. Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J.; Mordatch, I. Improving Factuality and Reasoning in Language Models through Multiagent Debate. *ArXiv* **2023**, *abs/2305.14325*. doi:10.48550/arXiv.2305.14325.

45. Luo, J.; Xiao, C.; Ma, F. Zero-Resource Hallucination Prevention for Large Language Models. *ArXiv* **2023**, *abs/2309.02654*. doi:10.48550/arXiv.2309.02654.

46. Creswell, A.; Shanahan, M. Faithful Reasoning Using Large Language Models. *ArXiv* **2022**, *abs/2208.14271*. doi:10.48550/arXiv.2208.14271.

47. Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q.; Zhou, D. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models, 2023, [arXiv:cs.CL/2201.11903].

48. Yeadon, W.; Halliday, D.P. Exploring durham university physics exams with large language models. *arXiv preprint arXiv:2306.15609* **2023**.

49. Singla, A. Evaluating ChatGPT and GPT-4 for Visual Programming. Proceedings of the 2023 ACM Conference on International Computing Education Research-Volume 2, 2023, pp. 14–15.

50. Zheng, C.; Liu, Z.; Xie, E.; Li, Z.; Li, Y. Progressive-Hint Prompting Improves Reasoning in Large Language Models. *ArXiv* **2023**, *abs/2304.09797*. doi:10.48550/arXiv.2304.09797.

51. Han, S.J.; Ransom, K.J.; Perfors, A.; Kemp, C. Inductive reasoning in humans and large language models. *Cognitive Systems Research* **2024**, *83*, 101155.

52. Liévin, V.; Hother, C.E.; Motzfeldt, A.G.; Winther, O. Can large language models reason about medical questions? *Patterns* **2023**.

53. Luo, L.; Lin, Z.; Liu, Y.; Shu, L.; Zhu, Y.; Shang, J.; Meng, L. Critique ability of large language models. *arXiv preprint arXiv:2310.04815* **2023**.

54. Feng, T.H.; Denny, P.; Wuensche, B.; Luxton-Reilly, A.; Hooper, S. More Than Meets the AI: Evaluating the performance of GPT-4 on Computer Graphics assessment questions. Proceedings of the 26th Australasian Computing Education Conference, 2024, pp. 182–191.

55. Bloom, B.S.; Engelhart, M.D.; Furst, E.J.; Hill, W.H.; Krathwohl, D.R.; others. *Taxonomy of educational objectives: The classification of educational goals. Handbook 1: Cognitive domain*; Longman New York, 1956.

56. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *Advances in neural information processing systems* **2017**, pp. 5998–6008.