

Using Large Language Models for education managements in Vietnamese with low resources

Duc Do Minh
UET-VNU
Vietnam National University
Hanoi, Vietnam
ducdm103.work@gmail.com

Vinh Nguyen Van
UET-VNU
Vietnam National University
Hanoi, Vietnam
vinhmv@vnu.edu.vn

Thang Dam Cong
Bac Ninh Teacher Training College
Bac Ninh, Vietnam
damcongthang@cbspbacninh.edu.vn

Abstract

Large language models (LLMs), such as GPT-4, Gemini 1.5, Claude 3.5 Sonnet, and Llama3, have demonstrated significant advancements in various NLP tasks since the release of ChatGPT in 2022. Despite their success, fine-tuning and deploying LLMs remain computationally expensive, especially in resource-constrained environments. In this paper, we proposed Viet-EduFrame, a framework specifically designed to apply LLMs to educational management tasks in Vietnamese institutions. Our key contribution includes the development of a tailored dataset, derived from student education documents at Hanoi VNU, which addresses the unique challenges faced by educational systems with limited resources. Through extensive experiments, we show that our approach outperforms existing methods in terms of accuracy and efficiency, offering a promising solution for improving educational management in under-resourced environments. While our framework leverages synthetic data to supplement real-world examples, we discuss potential limitations regarding broader applicability and robustness in future implementations.

1 Introduction

Most current tasks in Natural Language Processing (NLP) are dominated by large language models (LLMs) such as GPT4 and Gemini 1.5, which have set new benchmarks for performance. These models excel in a wide range of applications, demonstrating superior capabilities in understanding and generating human language.

In recent years, artificial intelligence (AI) and machine learning (ML) for education have received a great deal of interest and have been applied in various educational scenarios (Chen et al., 2020), (Xia et al., 2022), (Latif et al., 2023), (Denny et al., 2023), (Li et al., 2024). Educational data mining methods have been widely adopted in different aspects such as cognitive diagnosis (Batool et al.,

2022), knowledge tracking (Koedinger et al., 2015), and specifically question answering (Lende and Raghuvanshi, 2016), (Thiruvanantharajah et al., 2021), (Bhowmick et al., 2023).

Large language models (LLMs) have emerged as a powerful paradigm across different areas (Chen et al., 2023b), (Fan et al., 2023), (Jin et al., 2024), (Zeng et al., 2023), and have achieved state-of-the-art performances in multiple educational scenarios (Kasneji et al., 2023), (Li et al., 2023), (Yan et al., 2023). Existing work has found that LLMs can achieve student-level performance on standardized tests in a variety of subjects, including mathematics, physics, and computer science, on both multiple-choice and free-response problems. A recent study (Susnjak, 2022) reveals that ChatGPT is capable of generating logically consistent answers across disciplines, balancing both depth and breadth. Another quantitative analysis (Malinka et al., 2023) shows that students using ChatGPT (by keeping or refining the results from LLMs as their own answers) perform better than average students in some courses in the field of computer security.

Despite the global advancements, there remains a significant gap in the application of these technologies within the context of Vietnamese education, particularly in educational management. My research is among the first in Vietnam to explore these applications broadly in education and specifically in educational management. Due to the limitations of resources and data within Vietnamese institutions, this area has not yet received adequate attention. This scarcity of local studies and resources has driven us to undertake this research, aiming to bridge the gap and contribute to the growing body of knowledge in this critical field.

In this study, our main contributions can be summarized as follows:

- **Framework Proposal:** We propose a simple

yet highly effective framework for applying large language models (LLMs) to educational management tasks. This framework is designed to be easily implementable and adaptable within the constraints of Vietnamese educational institutions. To the best of our knowledge, this first study focuses applying LLMs for education in Vietnamese.

- **New Dataset:** We introduce a new dataset specifically tailored for educational management in Vietnam. This dataset addresses the unique challenges and characteristics of the Vietnamese educational context, providing a valuable resource for future research and development.
- **Model Development with Limited Resources:** We successfully develop and deploy a model using the limited computational resources available at our institution. This demonstrates the feasibility of implementing advanced AI solutions in resource-constrained environments and provides a blueprint for similar institutions.

2 Related work

2.1 Large language models in for study assisting

Providing students with timely learning support has been widely recognized as crucial in improving student engagement and learning efficiency during their independent studies (Dewhurst et al., 2000). Due to the limitation of prior algorithms in generating fixed-form responses, many of the existing study-assisting approaches face poor generalization challenges while being implemented in real-world scenarios (König et al., 2022). Fortunately, the appearance of LLMs brings revolutionary changes to this field. Using finetuned LLMs (Ouyang et al., 2022) to generate human-like responses, recent studies in LLM-based educational support have demonstrated promising results.

Contributing to the large-scale parameter size of LLMs and the enormous sized and diverse web corpus used during the pre-training phase, LLMs have been proven to be a powerful question zero-shot solver to questions spread from a wide spread of subjects, including math (Wu et al., 2023c) (Yuan et al., 2023), law (Bommarito and Katz, 2022) (Cui et al., 2023), medicine (Li'evin et al., 2022) (Thirunavukarasu et al., 2023), finance

(Wu et al., 2023b) (Yang et al., 2023), programming (Kazemitabaar et al., 2023) (avelka et al., 2023), language understands(Zhang et al., 2023). In addition, to further improve LLM's problem-solving performance while facing complicated questions, various studies have been actively proposed. For example, (Wei et al., 2022) proposes the Chain-of-Thought (CoT) prompting method, which guides LLMs to solve a challenging problem by decomposing it into simpler sequential steps. Other works exploit the strong in-context learning ability of LLMs and propose advanced few-shot demonstration-selection algorithms to improve LLM's problem-solving performance to general questions. (Chen et al., 2022) and (Gao et al., 2022b) leverage external programming tools to avoid calculation errors introduced during the textual problem-solving process of raw LLMs. (Wu et al., 2023a) regard chat-optimized LLMs as powerful agents and design a multi-agent conversation to solve those complicated questions through a collaborative process.

2.2 Education toolkit

Utilizing a chatbot powered by a Large Language Model (LLM) as an educational tool presents numerous benefits and opportunities. LLM chatbots can tailor their responses to meet the unique needs of each learner, offering personalized feedback and assistance. This ability to customize can cater to various learning styles, speeds, and preferences. They are available 24/7, making learning accessible at any time and from any place, which is especially advantageous for learners in different time zones or with diverse schedules. The interactive features of chatbots can make learning more engaging and enjoyable. They can mimic conversations, set up interactive learning scenarios, and give immediate feedback, which can be more effective than passive learning approaches. Chatbots can manage thousands of inquiries at once, providing a scalable solution for educational institutions to support a large number of students without needing more teaching staff. They can also automate repetitive teaching tasks, such as grading quizzes or offering basic feedback, enabling educators to concentrate on more complex and creative teaching duties. Notable examples of such chatbots include ChatGPT, Bing Chat, Google Bard, Perplexity, and Pi Pi.ai.

2.3 Textbook question answering

Textbook Question Answering (TQA) is a task that requires a system to comprehensively understand the multi-modal information from the textbook curriculum, spreading across text documents, images, and diagrams. The major challenge of textbook question answering is to comprehend the multi-modal domain-specific contexts as well as the questions, and then identify the key information to the questions.

Datasets (Kembhavi et al., 2017) presented the TQA dataset, designed to assess a system that integrates multi-modal contexts and a wide range of scientific topics. Comparable datasets, such as AI2D (Kembhavi et al., 2016), DVQA (Kafle et al., 2018), and VLQA (Sampat et al., 2020), have been developed to facilitate research in multi-modal reasoning within the scientific domain. Nonetheless, these datasets lack annotated explanations for answers in the form of supporting facts. SCIENCEQA (Lu et al., 2022) is a comprehensive textbook question-answering dataset that includes annotated lectures and explanations. This dataset is derived from elementary and high school science curricula, covering a variety of science topics such as natural science, social science, and language science. Recently, the TheoremQA dataset has been released, which includes textbook questions at the university level (Chen et al., 2023a). Beyond the scientific domain, there are datasets focused on the medical field. MEDQA (Jin et al., 2020) and MedMCQA (Pal et al., 2022) are two medical question-answering datasets that encompass a broad range of healthcare topics, derived from both real-world scenarios and simulated exams.

Methods. From a technical perspective, textbook question answering is inherently similar to visual question answering (VQA) (Dosovitskiy et al., 2020), (Gao et al., 2018), (Gao et al., 2022a). Traditional VQA approaches use RNNs to encode the question and CNNs to encode the image (Agrawal et al., 2015), (Malinowski et al., 2015). The multi-modal information is then fused to understand the questions. Additionally, other methods that utilize spatial attention (Lu et al., 2016), (Noh and Han, 2016), (Xu et al., 2015), (Yang et al., 2015), compositional strategies (Andreas et al., 2016), and bilinear pooling schemes (Fukui et al., 2016), (Liu et al., 2022) have been proposed to enhance VQA performance.

While VQA and textbook question answering share significant similarities, textbook question answering requires domain-specific knowledge for the accompanying context and innovative integration of diagrams and tables. To address this gap, (Ram et al., 2021) proposed a pre-training schema tailored for question answering. Specifically, their method improves performance in textbook question answering by masking recurring span selections and selecting the correct span in the passage, even when only a hundred examples are available in specific domains. An adversarial training framework is also adapted for domain generalization (Lee et al., 2019), enabling question-answering models to learn domain-invariant features. (Xu et al., 2022) introduced a novel Pre-trained Machine Reader as an enhancement of pre-trained Masked Language Models (MLMs), which addresses the discrepancy between model pre-training and downstream fine-tuning for specific domain MLMs. To comprehend diagrams and tables, graph-based parsing methods have been developed to extract concepts from diagrams (Kembhavi et al., 2016) by converting a diagram into a diagram parse graph. Optical Character Recognition (OCR) is employed to identify chart-specific answers from the charts, which are then aligned with the questions (Poco and Heer, 2017), (Kafle et al., 2018).

Our research is different from previous works in some significant ways:

- **First**, we have developed a simple yet effective framework for the textbook question-answering problem. This framework has proven to be both efficient and robust, delivering high performance within a short development cycle.
- **Second**, leveraging this framework, we have created a dedicated dataset specifically tailored for the training management process at the Vietnam National University of Hanoi. This dataset is instrumental in enhancing the quality and effectiveness of training management, marking a substantial contribution to the educational resources available for Vietnamese institutions.

3 Dataset

The use of data in the field of educational management presents several significant challenges, particularly when developing a question-answering

system for the Vietnamese language. These challenges include:

- **Institutional Variability:** Each educational institution must comply with the regulations set forth by the Ministry of Education. However, beyond these mandatory guidelines, institutions often have additional rules and policies specific to their own organization or the larger entity they are affiliated with. This variability can lead to inconsistencies in data structure, terminology, and reporting practices, complicating the task of creating a unified dataset.
- **Data Standardization:** Due to the diverse regulatory requirements and internal policies across different institutions, standardizing data becomes a complex process. Ensuring consistency and compatibility of data from various sources is essential for effective analysis and model training but is difficult to achieve given the heterogeneity of the data.
- **Data Availability and Quality:** As one of the first studies addressing the question-answering problem in the Vietnamese language within the educational management domain, there is a scarcity of readily available datasets. Existing datasets in other languages or educational contexts may not be directly applicable due to linguistic and contextual differences. Therefore, sourcing high-quality data externally is challenging, necessitating the creation of a new dataset from scratch.
- **Data Collection and Annotation:** Building a new dataset requires significant effort in data collection and annotation. This process involves gathering data from various educational institutions, ensuring its accuracy and relevance, and annotating it to create a structured dataset suitable for training machine learning models. The annotation process, in particular, is time-consuming and demands a deep understanding of the educational domain.

Addressing these challenges is crucial for the success of our research. By acknowledging and systematically tackling these issues, we aim to build a robust and reliable dataset that will facilitate the development of effective AI solutions for educational management in Vietnam.

3.1 Building data

In this subsection, I describe the process of constructing a dataset from the "Regulations on Student Affairs of Vietnam National University"(VNU) to train a model for question-answering tasks. By using prompts, we generate data points that each consist of a "context," "question," and "answer." This structured approach ensures comprehensive coverage of the regulations and facilitates the creation of a robust dataset for training. The process consists of five critical steps: data preprocessing, data analysis and prompt design, data generation using prompts and LLMs, and data quality evaluation.

The first step **data preprocessing** involves preparing the raw data for subsequent analysis and prompt generation. This includes:

- **Data Cleaning:** Removing any irrelevant information, and duplicates, and ensuring consistency in formatting.
- **Text Segmentation:** Breaking down the regulations into manageable sections that can be used as context for generating questions and answers.
- **Whitespace and Extraneous Character Removal:** Removing unnecessary spaces and characters to ensure clean text.
- **Spell Checking:** Correcting any spelling errors in the text.
- **Math Formula Conversion:** Converting mathematical formulas into KATEX format for consistent representation.

After preprocessing the data, the next step is to **analyze the content and design effective prompts**. This involves:

- **Content Analysis:** Identifying key themes, rules, and guidelines within the regulations.
- **Prompt Crafting:** Developing specific prompts that will be used to generate questions and answers. Each prompt focuses on different aspects of the regulations, ensuring comprehensive coverage.
- **Using technique prompting** Chain of Thought, Self-Consistency Chain of Thought, and Tree of Thought: Employing advanced prompting techniques to enhance the generation process.

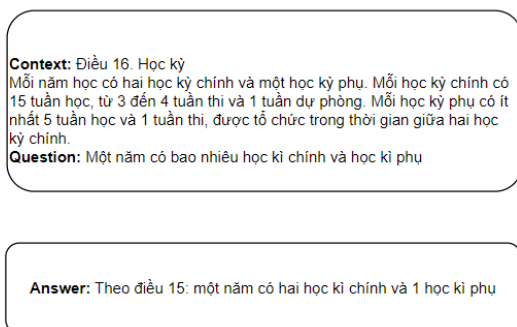
The next steps are **data generation using prompts and LLMs, and data quality evaluation**. To generate the desired dataset, we utilized prompts that were meticulously designed in the previous phase. These prompts were fed into large language models (LLMs) such as GPT-3,5 turbo, which then generated a comprehensive set of synthetic data. The generation process was systematic and aimed to produce data that closely aligned with our research objectives and covered the necessary range of scenarios.

The quality of the generated data was evaluated using both automated metrics and human assessment. Specifically, we employed ROUGE and BLEU scores to quantify the relevance and coherence of the generated text. These metrics provided an objective measure of how well the generated data matched the expected output in terms of n-gram overlap and sequence similarity.

In addition to automated metrics, human evaluators conducted a qualitative review of the generated data. These domain experts assessed the data for relevance, coherence, and diversity, ensuring that the synthetic data met the high standards required for our study. This dual approach of combining quantitative scores with qualitative human judgment ensured a robust evaluation of the generated dataset, confirming its suitability for subsequent analyses and experiments.

Here is an example of the dataset

Figure 1: The examples of Question Answering in the education domain



4 Methodology

In this section, we detail the methodology employed to address the question-answering problem within the domain of university educational management in 2. Our approach encompasses several

key stages: leveraging a Large Language Model (LLM) for initial data pre-labeling, human labeling for data refinement, training the model, evaluating its performance, and conducting a thorough analysis of the results. Each step in this pipeline is meticulously designed to ensure accuracy and effectiveness, tailored to the specific needs and constraints of the educational context in Vietnam.

We will systematically describe each stage of our methodology as follows:

- **Large Language Model (LLM):**An overview of the LLM utilized in our study, highlighting its features and advantages in handling natural language processing tasks.
- **Pre-labeling:** A description of the pre-labeling process using the LLM to provide initial annotations for the dataset, which sets a foundation for further refinement.
- **Human Labeling:** An explanation of the human labeling process, emphasizing its role in ensuring high-quality data by correcting and improving the initial LLM-generated labels.
- **Training:** Details on the training phase, including the algorithms and techniques applied to build a robust question-answering model.
- **Evaluation:** Presentation of the evaluation methods and criteria used to assess the model's performance, ensuring it meets the desired standards of accuracy and reliability.
- **Analysis:**A comprehensive analysis of the results obtained from the evaluation, providing insights into the model's strengths and areas for improvement.

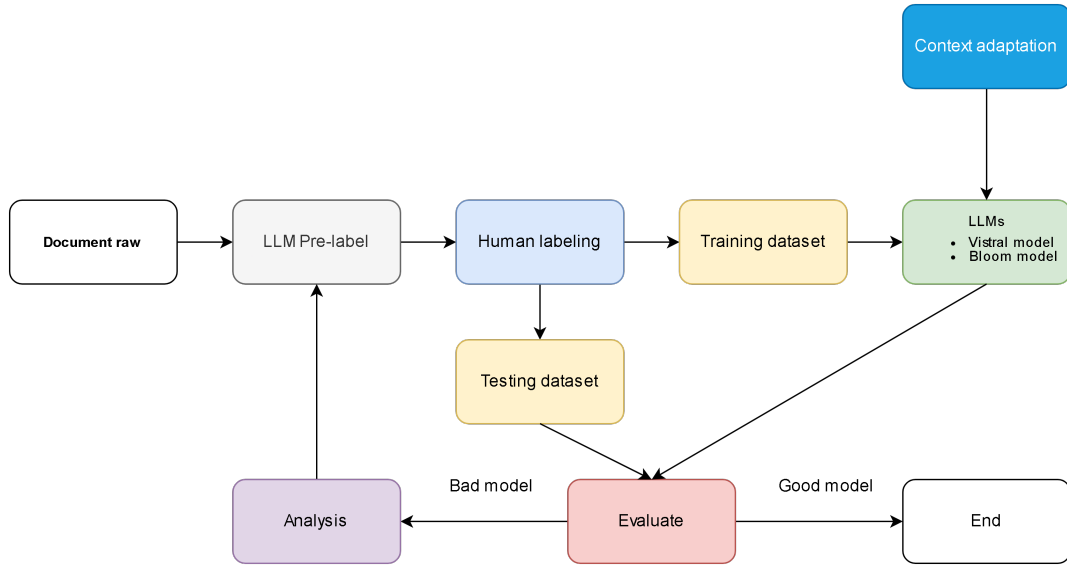
4.1 Pre labeling and human labeling

With two steps using LLMs pre-labeling và human labeling, I illustrated in section 3 building data.

4.2 Training and context adaptation

In this subsection, we describe the training process and context adaptation techniques employed to enhance the question-answering capabilities of our model, particularly tailored to university educational management.

Figure 2: Overview of our framework



4.2.1 Training

Model Vistral

Vistral (Vo, 2024) is a deep learning model that uses many transformer decoder layers to generate coherent and natural language text. The model was pre-trained on a large corpus of text data using an unsupervised learning approach, which enabled it to learn the statistical patterns and structures of natural language. Vistral has been widely used for various NLP tasks such as language translation, question-answering, text summarization, and even creative writing. As of now April 2024, the Vistral model is the highest-scoring public model on the VMLU leaderboard. Vistral model is an innovative Large Language Model designed expressly for the Vietnamese language.

- Rolling Buffer Cache
- Sliding-Window Attention
- Pre-fill and Chunking

Sliding Window Attention utilizes the multiple layers of a transformer to access information beyond a defined window size W . In this method, the hidden state at position i in layer k , denoted as h_i , attends to all hidden states in the preceding layer within the range from $i - W$ to i . This process allows h_i to recursively access tokens from the input layer at a distance of up to $W \times k$ tokens.

Rolling Buffer Cache. By having a fixed attention span, we can manage our cache size with a rolling buffer cache. This cache has a set size of

W , and the keys and values for timestep i are saved in the position $i \bmod W$ of the cache. Consequently, when position i exceeds W , the older values in the cache are overwritten, preventing the cache size from growing indefinitely.

Pre-fill and Chunking. When generating a sequence, tokens must be predicted one at a time, as each token depends on the previous ones. However, since the prompt is known beforehand, we can pre-fill the (k, v) cache with the prompt. If the prompt is very large, it can be divided into smaller chunks, and the cache can be pre-filled with these chunks. The window size can be used as the chunk size. For each chunk, it is necessary to compute the attention over both the cache and the chunk.

Model Bloom BLOOM is a powerful autoregressive Large Language Model (LLM) designed to extend text from a given prompt, utilizing extensive computational resources on massive text datasets. This capability allows it to produce fluent text in 46 different languages and 13 programming languages, making it almost indistinguishable from human-written content. Additionally, BLOOM can be directed to undertake text-related tasks it wasn't specifically trained for by framing them as text-generation problems.

Modeling Details Several key innovations were incorporated into the BLOOM model to enhance its performance and stability:

ALiBi Positional Embeddings: The model employs ALiBi (Attention Linear Bias) positional embeddings instead of traditional positional embed-

dings. ALiBi attenuates attention scores based on the distance between keys and queries, which results in smoother training dynamics and improved performance.

Embedding LayerNorm: An additional layer normalization step is applied immediately after the embedding layer. This modification was implemented to improve training stability, especially considering the use of bfloat16 precision in the final training phase, which offers more stability than float16.

Low rank Adaptation

For a given pretrained weight matrix $W_0 \in \mathbb{R}^{d \times k}$, LoRA introduces two trainable weight matrices, $W_{up} \in \mathbb{R}^{d \times r}$ and $W_{down} \in \mathbb{R}^{r \times k}$ where the rank $r \ll \min(d, k)$, operating in parallel to W_0 . Let h_{in} represent the input. Under normal conditions, the output through W_0 is $h_{out} = W_0 h_{in}$. Instead, LoRA modifies this output by introducing an incremental update ΔW that encapsulates task-specific knowledge:

$$h_{out} = W_0 h_{in} + \frac{\alpha}{r} \Delta W h_{in} = W_0 h_{in} + \frac{\alpha}{r} W_{up} W_{down} h_{in} \quad (1)$$

where α denotes a scaling factor. At the onset of training, W_{down} is initialized using a random Gaussian distribution, while W_{up} is initialized to zero, ensuring that ΔW initially holds a value of zero. LoRA is straightforward to implement and has been evaluated on models with up to 175 billion parameters. In this research, I use this method for the model Bloom and Vistral-7B. Once fine-tuning is complete, LoRA’s adaptive weights seamlessly integrate with the pre-trained backbone weights. This integration ensures that LoRA maintains the model’s efficiency, adding no extra burden during inference. The number of parameters training is reduced $dk/(d+k)/r$ times.

4.2.2 Context Adaptation

Context adaptation is crucial for activating the model’s question-answering capabilities. We enhance the training data by incorporating detailed instructions and contextual cues that guide the model in understanding and generating accurate responses to educational queries.

By adding specific instructions, we provide the model with explicit examples of how to approach different types of questions within the educational domain. These instructions act as triggers, enabling

the model to apply its learned knowledge effectively and respond accurately to complex queries.

Our training and context adaptation approach ensures that the models are not only finely tuned to our dataset but also contextually aware, enhancing their ability to provide precise and relevant answers in the context of university educational management. The combination of dual-model training and LoRA, along with detailed context adaptation, significantly boosts the model’s performance and usability in real-world applications.

4.3 Evaluate

Exact Match (EM): For each question-answer pair, if the characters of the MRC system’s predicted answer exactly match the characters of (one of) the gold standard answer(s), $EM = 1$, otherwise $EM = 0$. EM is a stringent all-or-nothing metric, with a score of 0 for being off by a single character. When evaluating against a negative question, if the system predicts any textual span as an answer, it automatically obtains a zero score for that question.

F1-score: F1-score is a popular metric for natural language processing and is also used in machine reading comprehension. F1-score is estimated over the individual tokens in the predicted answer against those in the gold standard answers. The F1-score is based on the number of matched tokens between the predicted and gold standard answers.

$$\text{Precision} = \frac{\text{the number of matched tokens}}{\text{the total tokens in the predicted answer}} \quad (2)$$

$$\text{Recall} = \frac{\text{the number of matched tokens}}{\text{the total tokens in the gold standard answer}} \quad (3)$$

$$\text{F1-score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4)$$

5 Result and Experiment

5.1 Statistic of dataset

In this subsection, we present a comprehensive statistical analysis of our dataset, which includes an in-depth survey of the lengths and averages of contexts, questions, and answers. Understanding these metrics is crucial for evaluating the overall quality and characteristics of the data used in our experiments.

5.2 Data review

In our study, we categorize the dataset into five distinct levels of question-answering data quality:

Table 1: Statistic of dataset

| | context length | question length | answer length |
|-------|----------------|-----------------|---------------|
| count | 985.00 | 985.00 | 985.00 |
| mean | 882.48 | 74.03 | 415.60 |
| std | 742.12 | 32.59 | 342.66 |
| min | 49.00 | 15.00 | 21.00 |
| 25% | 324.00 | 54.00 | 166.00 |
| 50% | 611.00 | 71.00 | 298.00 |
| 75% | 1371.00 | 86.00 | 569.00 |
| max | 4446.00 | 289.00 | 2163.00 |

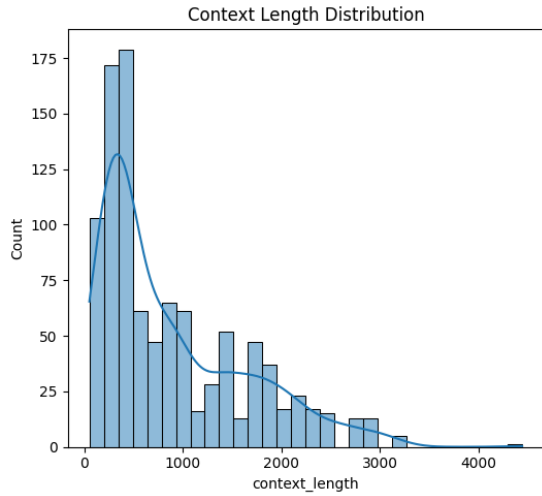


Figure 3: Context Length Distribution

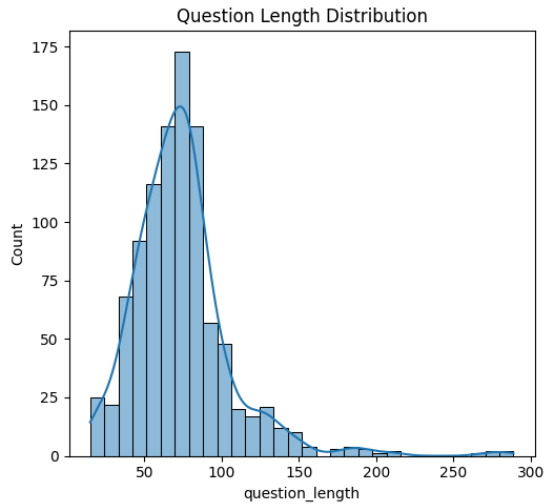


Figure 4: Question Length Distribution

Very Good, Good, Medium, Bad, and Very Bad. These levels are comprehensively described in Table 3

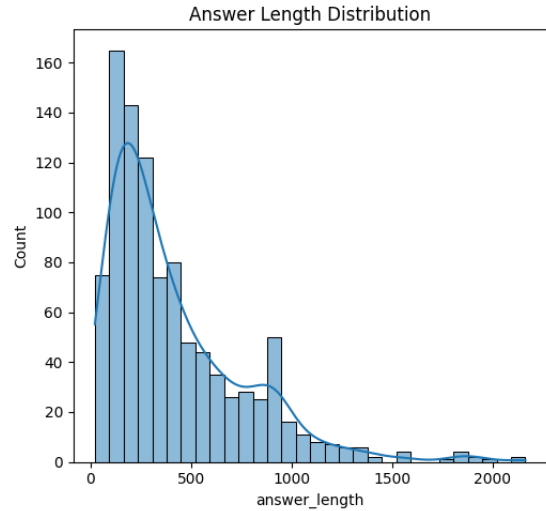


Figure 5: Answer Length Distribution

Table 2: Levels of Data Quality in Question Answering

| Type | Description |
|-----------|--|
| Very good | Answers at this level are completely accurate and directly address the question posed. They exhibit a perfect understanding of the query and provide comprehensive, precise information. The content is well-structured and leaves no room for ambiguity. |
| Good | Answers in this category are mostly accurate and address the main aspects of the question. They may lack some minor details or have slight imprecisions but still provide a reliable and useful response. These answers are generally clear and relevant. |
| Medium | Answers at this level are somewhat accurate but may be incomplete or partially incorrect. They provide relevant information but may miss key details or present some minor inaccuracies. The response could be clearer or more comprehensive. |
| Bad | Answers in this category are largely inaccurate or irrelevant. They may partially address the question but contain significant errors or omissions. The response may be confusing, vague, or off-topic, requiring substantial correction or clarification. |
| Very Bad | Answers at this level are completely incorrect or irrelevant. They fail to address the question in any meaningful way, providing no useful information. The response might be entirely off-topic or nonsensical, reflecting a fundamental misunderstanding of the query. |

5.3 Result of model

In this section, we present the performance of the Bloom and Vistral models. The results are evaluated using the training and validation loss metrics, as well as a comparison of the exact match (Exact)

Table 3: Percentage Data Quality

| Type | Number | Percentage |
|--------------|-------------|--------------|
| Very good | 631 | 54.92 % |
| Good | 325 | 28.28 % |
| Medium | 103 | 8.96 % |
| Bad | 78 | 6.78 % |
| Very Bad | 12 | 1.05 % |
| Total | 1149 | 100 % |

and F1 scores.

I implement hyperparameters with full fine-tuning model in table 5 and hyperparameter using LoRA for tuning model in table 5.

Table 4: Hyperparameter of Bloom and Vistral models

| Model | Bloom | Vistral |
|--------------|-------|---------|
| β_1 | 0.9 | 0.9 |
| β_2 | 0.999 | 0.999 |
| warmup ratio | 0.05 | 0.05 |
| weight decay | 0.01 | 0.01 |
| batch size | 8 | 4 |
| max length | 1024 | 1024 |
| num epochs | 10 | 10 |

Table 5: Hyperparameter of Bloom and Vistral models with LoRA

| Model | Bloom | Bloom |
|--------------|-------|-------|
| β_1 | 0.9 | 0.9 |
| β_2 | 0.999 | 0.999 |
| warmup ratio | 0.05 | 0.05 |
| weight decay | 0.01 | 0.01 |
| batch size | 4 | 8 |
| max length | 1024 | 1024 |
| num epochs | 10 | 10 |
| Rank LoRA | 128 | 128 |
| LoRA dropout | 0.1 | 0.1 |

Training and Validation Loss Bloom Model:

The training and validation loss curves for the Bloom model are shown in figures ?? and 7, respectively. Additionally, the training loss of Bloom model and LoRA method have training loss in figure 8 and validation loss illustrated in figure 9.

Vistral Model:

Similarly, the training and validation loss curves for the Vistral model are depicted in Figures 10 and 11. The Vistral model shows a rapid decrease in training loss, and the validation loss also reduces steadily, demonstrating good generalization performance. Furthermore, in figure 12, 13 present loss

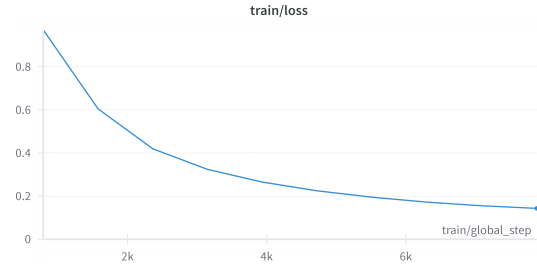


Figure 6: Training Loss of Bloom Model

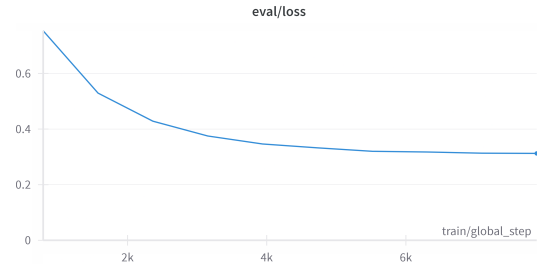


Figure 7: Validation Loss of Bloom Model

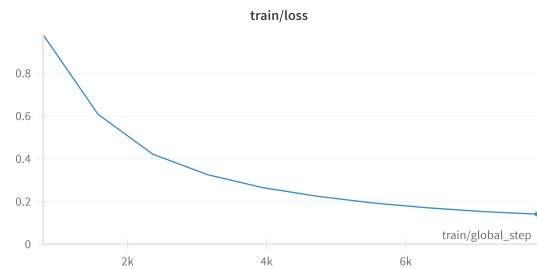


Figure 8: Training Loss of Bloom Model

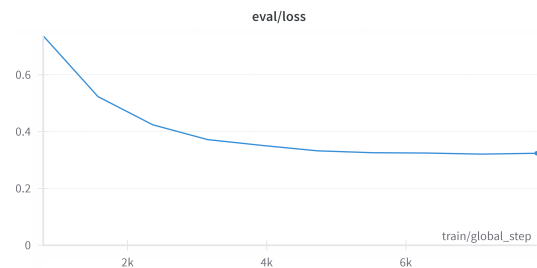


Figure 9: Validation Loss of Bloom Model + LoRA

of training and validate phrases respectively.

Comparison of Bloom and Vistral Models Table 6 provides a comparison of the Exact and F1 scores for both the Bloom and Vistral models. The Vistral model outperforms the Bloom model in both metrics, indicating its superior performance in terms of both accuracy and the quality of predictions.

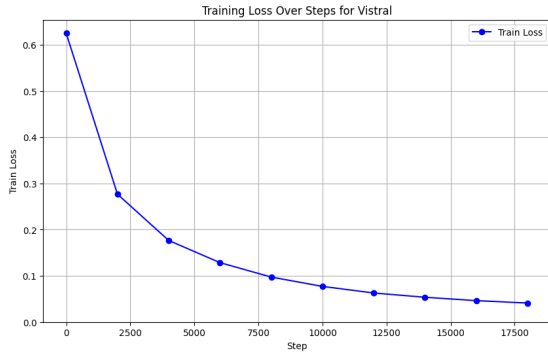


Figure 10: Training Loss of Vistral Model

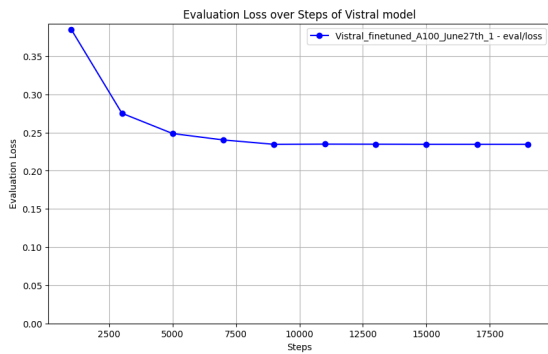


Figure 11: Validation Loss of Vistral Model

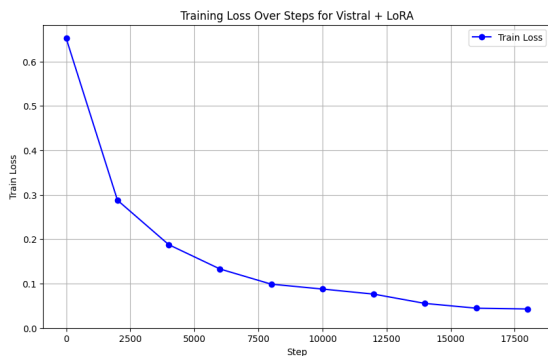


Figure 12: Training Loss of Vistral Model + LoRA

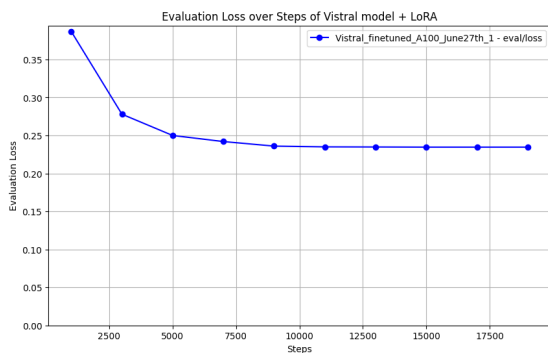


Figure 13: Validation Loss of Vistral Model + LoRA

Table 6: Overall result

| Metric | Exact | F1-score |
|---------------------------|-------|----------|
| Bloom model + LoRA | 33.89 | 72.36 |
| Vistral + LoRA | 43.23 | 81.24 |
| Bloom model | 34.23 | 73.16 |
| Vistral model | 43.72 | 81.57 |

Table 7: Resource usage of language models

| Model | Time training per epoch | Ram-GPU used |
|---------------------------|-------------------------|--------------|
| Bloom model + LoRA | 1.5 hours | 16 GB |
| Vistral + LoRA | 6 hours | 32 GB |
| Bloom model | 5 hours | 29 GB |
| Vistral model | 14 hours | 61.2 GB |

6 Analysis and discussion

6.1 Performance Metrics

- **Bloom model + LoRA vs. Bloom model:** The Bloom model with LoRA shows a slight decrease in Exact and F1-score compared to the Bloom model without LoRA. The Exact score drops from 34.23 to 33.89, and the F1 score decreases from 73.16 to 72.36. This suggests that LoRA might slightly affect the performance of the Bloom model in terms of these metrics.
- **Vistral + LoRA vs. Vistral:** The Vistral model with LoRA also exhibits a minor reduction in performance compared to the Vistral model without LoRA. The Exact score drops from 43.72 to 43.23, and the F1 score decreases from 81.57 to 81.24. This indicates that the inclusion of LoRA may have a small impact on the Vistral model's performance.
- **Bloom model vs. Vistral:** Comparing the two models, Vistral consistently outperforms Bloom in both Exact and F1-score, both with and without LoRA. This demonstrates that the Vistral model is more effective in capturing and processing the information needed for higher precision and overall accuracy.

6.2 Resource Utilization

- **Training Time:** The training time per epoch is significantly lower for models using LoRA. The Bloom model with LoRA takes 1.5 hours per epoch, whereas without LoRA, it takes 5 hours. Similarly, the Vistral model with LoRA takes 6 hours per epoch, compared to 14 hours without LoRA. This reduction in training time highlights the efficiency of the LoRA method in speeding up the training process.
- **GPU RAM Usage:** Models with LoRA also require less GPU RAM. The Bloom model with LoRA uses 16 GB, while the original Bloom model uses 29 GB. The Vistral model with LoRA uses 32 GB, compared to 61.2 GB for the Vistral model without LoRA. This reduction in memory usage indicates that LoRA helps in optimizing resource utilization during training.

6.3 Real-world inference

Example in table 8 there are four reasons explain why this is a good answer:

- **Comprehensive and Detailed:** The answer includes the main responsibilities of the lecturers such as imparting scientific ambition, a passion for learning, scientific research, scientific thinking, and creative ability. These elements are clearly stated in the context provided.
- **Clear and Understandable:** The answer is articulated clearly and understandably, making it easy for the reader to grasp the responsibilities of the lecturers.
- **Contextual Connection:** The answer is closely linked to the context provided, ensuring that the information presented is accurate and relevant.
- **Repetition for Emphasis:** Repeating the key points at the end of the answer emphasizes the lecturers' responsibilities and highlights the main aspects that need to be noted.

In table 9 illustrates the bad answer. This poor answer does not provide a comprehensive response based on the provided context. It lacks details and does not address the specific aspects mentioned in the context, such as training programs, study time,

Table 8: Good answer in model

| | |
|-----------------|---|
| Context | Điều 33. Giảng viên c) Ngoài những nghĩa vụ chung đối với giảng viên , các giảng viên tham gia đào tạo tài năng , chất lượng cao còn có trách nhiệm sau : - Truyền thụ và tạo lập cho sinh viên hoài bão khoa học , niềm say mê học tập , nghiên cứu khoa học , tư duy khoa học và năng lực sáng tạo , động viên ,khuyến khích sinh viên rèn luyện phẩm chất chính trị, đạo đức tác phong của người cán bộ nghiên cứu khoa học ; - Tự giác áp dụng và hướng dẫn sinh viên áp dụng phương pháp dạy -học , kiểm tra đánh giá tiên tiến phù hợp ; - Hướng dẫn sinh viên nghiên cứu khoa học ; Cung cấp học liệu cho sinh viên trước giờ lên lớp ở một trong các dạng sau : giáo trình ;bài giảng ; đề cương chi tiết bài giảng cùng danh mục tài liệu tham khảo có trong thư viện . Nội dung bài giảng phải cơ bản , hiện đại, trình bày được các thành tựu nghiên cứu mới trong nước và quốc tế, phù hợp với mục tiêu đào tạo , giới thiệu các tài liệu tham khảo phong phú cho sinh viên. " |
| Question | Giảng viên có trách nhiệm gì đối với sinh viên trong quá trình đào tạo? |
| Answer | Giảng viên có trách nhiệm truyền thụ và tạo lập cho sinh viên hoài bão khoa học, niềm say mê học tập, nghiên cứu khoa học, tư duy khoa học và năng lực sáng tạo. Hãy suy nghĩ từng bước. Vì vậy, câu trả lời là giảng viên có trách nhiệm truyền thụ và tạo lập cho sinh viên hoài bão khoa học, niềm say mê học tập, nghiên cứu khoa học, tư duy khoa học và năng lực sáng tạo. |

organization and management of training, scientific research, rights and obligations of lecturers, advisors, students, examinations, assessments, and graduation recognition.

Table 9: Bad answer in model

| | |
|-----------------|--|
| Context | Điều 1. Phạm vi điều chỉnh và đối tượng áp dụng 1. Quy chế này quy định về đào tạo đại học theo hình thức chính quy và hình thức vừa làm vừa học tại Đại học Quốc gia Hà Nội ,bao gồm :chương trình đào tạo và thời gian học tập ;tổ chức và quản lý đào tạo; nghiên cứu khoa học của sinh viên ; quyền lợi và nghĩa vụ của giảng viên , cố vấn học tập, sinh viên ; kiểm tra, thi và đánh giá kết quả học tập ;công nhận tốt nghiệp |
| Question | Quy chế đào tạo đại học quy định về những gì? |
| Answer | Quy chế đào tạo đại học quy định về đào tạo. |

Conclusion and Limitations

6.4 Conclusion

In this paper, we present a simple and effective framework for applying large language models (LLMs) to educational domain. We conduct the experiments with fine-tuning methods on resource-constrained environments to optimally leverage existing GPU capabilities and hardware. Our results demonstrated that using LLMs models for vietnamese improved performance by over 10 points compared to previous model. This significant improvement highlights the effectiveness of our approach in maximizing the potential of limited computational resources.

6.5 Limitations

In this study and in the realm of natural language processing, particularly in the application of question-answering (QA) systems for educational management in Vietnamese, several limitations of current models and data quality have been identified. These limitations are crucial to understand for the continued development and improvement of such systems.

1. Reasoning Capabilities of the Model

- **Logical Reasoning:** The models may produce answers that lack coherent logical structure or fail to follow a clear line of reasoning, especially for complex or multi-step problems.

- **Contextual Understanding:** While models can understand the context to a certain extent, they often miss subtle nuances and deeper connections within the provided context, leading to less accurate or irrelevant responses.

2. Contextual Errors and Ambiguity

- **Error in Capturing Context:** Models sometimes fail to capture the full context of a question, particularly when the context is lengthy or contains intricate details.
- **Ambiguity in Responses:** Due to the models' probabilistic nature, they can produce responses that are ambiguous or vague, which can be particularly problematic in educational management where precision is crucial.

3. Lack of Specialized Knowledge

- **Handling Specific Regulations:** The models might not fully grasp the specific regulations and guidelines unique to different educational institutions or contexts, leading to incorrect or incomplete answers.
- **Domain-Specific Expertise:** The absence of deep domain expertise means that the models might misinterpret or overlook critical aspects of educational management tasks.

References

- Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C. Lawrence Zitnick, Devi Parikh, and Dhruv Batra. 2015. [Vqa: Visual question answering](#). *International Journal of Computer Vision*, 123:4 – 31.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. [Learning to compose neural networks for question answering](#). *ArXiv*, abs/1601.01705.
- Jaromr avelka, Arav Agarwal, Marshall An, Christopher Bogart, and Majd F. Sakr. 2023. [Thrilled by your progress! large language models \(gpt-4\) no longer struggle to pass assessments in higher education programming courses](#). *Proceedings of the 2023 ACM Conference on International Computing Education Research - Volume 1*.

- Saba Batool, Junaid Rashid, Muhammad Wasif Nisar, Jungeun Kim, Hyuk-Yoon Kwon, and Amir Hussain. 2022. [Educational data mining to predict students' academic performance: A survey study](#). *Education and Information Technologies*, 28:905–971.
- Ayan Kumar Bhowmick, Ashish Jagmohan, Aditya Vempaty, Prasenjit Dey, Leigh Ann Edwards Hall, Jeremy Hartman, Ravi Kokku, and Hema Maheshwari. 2023. [Automating question generation from educational text](#). In *SGAI Conferences*.
- Michael James Bommarito and Daniel Martin Katz. 2022. [Gpt takes the bar exam](#). *ArXiv*, abs/2212.14402.
- Lijia Chen, Pingping Chen, and Zhijian Lin. 2020. [Artificial intelligence in education: A review](#). *IEEE Access*, 8:75264–75278.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2022. [Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks](#). *ArXiv*, abs/2211.12588.
- Wenhu Chen, Ming Yin, Max W.F. Ku, Yixin Wan, Xueguang Ma, Jianyu Xu, Tony Xia, Xinyi Wang, and Pan Lu. 2023a. [Theoremqa: A theorem-driven question answering dataset](#). *ArXiv*, abs/2305.12524.
- Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Haifang Wen, Xiaochi Wei, Shuaiqiang Wang, Dawei Yin, Wenqi Fan, Hui Liu, and Jiliang Tang. 2023b. [Exploring the potential of large language models \(llms\) in learning on graphs](#). *ACM SIGKDD Explorations Newsletter*, 25:42 – 61.
- Jiaxi Cui, Zongjia Li, Yang Yan, Bohua Chen, and Li Yuan. 2023. [Chatlaw: A multi-agent collaborative legal assistant with knowledge graph enhanced mixture-of-experts large language model](#).
- Paul Denny, James Prather, Brett A. Becker, James Finnie-Ansley, Arto Hellas, Juho Leinonen, Andrew Luxton-Reilly, Brent N. Reeves, Eddie Antonio Santos, and Sami Sarsa. 2023. [Computing education in the era of generative ai](#). *Communications of the ACM*, 67:56 – 67.
- David Dewhurst, Hamish Macleod, and Tracey A. M. Norris. 2000. [Independent student learning aided by computers: an acceptable alternative to lectures?](#) *Comput. Educ.*, 35:223–241.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2020. [An image is worth 16x16 words: Transformers for image recognition at scale](#). *ArXiv*, abs/2010.11929.
- Wenqi Fan, Zihuai Zhao, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Jiliang Tang, and Qing Li. 2023. [Recommender systems in the era of large language models \(llms\)](#). *ArXiv*, abs/2307.02046.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. [Multimodal compact bilinear pooling for visual question answering and visual grounding](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Feng Gao, Q. Ping, Govind Thattai, Aishwarya N. Reganti, Yingting Wu, and Premkumar Nataraajan. 2022a. [Transform-retrieve-generate: Natural language-centric outside-knowledge visual question answering](#). *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5057–5067.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2022b. [Pal: Program-aided language models](#). *ArXiv*, abs/2211.10435.
- Peng Gao, Zhengkai Jiang, Haoxuan You, Pan Lu, Steven C. H. Hoi, Xiaogang Wang, and Hongsheng Li. 2018. [Dynamic fusion with intra- and inter-modality attention flow for visual question answering](#). *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6632–6641.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2020. [What disease does this patient have? a large-scale open domain question answering dataset from medical exams](#). *ArXiv*, abs/2009.13081.
- Ming Jin, Yifan Zhang, Wei Chen, Kexin Zhang, Yuxuan Liang, Bin Yang, Jindong Wang, Shirui Pan, and Qingsong Wen. 2024. [Position: What can large language models tell us about time series analysis](#).
- Kushal Kafle, Scott D. Cohen, Brian L. Price, and Christopher Kanan. 2018. [Dvqa: Understanding data visualizations via question answering](#). *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5648–5656.
- Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, George Louis Groh, Stephan Günnemann, Eyke Hüllermeier, Stephan Krusche, Gitta Kutyniok, Tilman Michaeli, Claudia Nerdel, Jürgen Pfeffer, Oleksandra Poquet, Michael Sailer, Albrecht Schmidt, Tina Seidel, Matthias Stadler, Jochen Weller, Jochen Kuhn, and Gjergji Kasneci. 2023. [Chatgpt for good? on opportunities and challenges of large language models for education](#). *Learning and Individual Differences*.
- Majeed Kazemitabaar, Xinying Hou, Austin Henley, Barb Ericson, David Weintrop, and Tovi Grossman. 2023. [How novices use llm-based code generators to solve cs1 coding tasks in a self-paced learning environment](#). *Proceedings of the 23rd Koli Calling International Conference on Computing Education Research*.

- Aniruddha Kembhavi, Michael Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. [A diagram is worth a dozen images](#). *ArXiv*, abs/1603.07396.
- Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension](#). *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5376–5384.
- K. Koedinger, Sidney K. D’Mello, Elizabeth Mclaughlin, Zachary A. Pardos, and Carolyn Penstein Rosé. 2015. [Data mining and education](#). *Wiley interdisciplinary reviews. Cognitive science*, 6 4:333–353.
- Claudia M. König, Christin Karrenbauer, and Michael H. Breitner. 2022. [Critical success factors and challenges for individual digital study assistants in higher education: A mixed methods analysis](#). *Education and Information Technologies*, 28:4475 – 4503.
- Ehsan Latif, Gengchen Mai, Matthew Nyaaba, Xuansheng Wu, Ninghao Liu, Guoyu Lu, Sheng Li, Tianming Liu, and Xiaoming Zhai. 2023. [Artificial general intelligence \(agi\) for education](#). *ArXiv*, abs/2304.12479.
- Seanie Lee, Donggyu Kim, and Jangwon Park. 2019. [Domain-agnostic question-answering with adversarial training](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Sweta P. Lende and Mukesh M. Raghuwanshi. 2016. [Question answering system on education acts using nlp techniques](#). *2016 World Conference on Futuristic Trends in Research and Innovation for Social Welfare (Startup Conclave)*, pages 1–6.
- Hang Li, Tianlong Xu, Chaoli Zhang, Eason Chen, Jing Liang, Xing Fan, Haoyang Li, Jiliang Tang, and Qingsong Wen. 2024. [Bringing generative ai to adaptive learning in education](#). *ArXiv*, abs/2402.14601.
- Qingyao Li, Lingyue Fu, Weiming Zhang, Xianyu Chen, Jingwei Yu, Wei Xia, Weinan Zhang, Ruiming Tang, and Yong Yu. 2023. [Adapting large language models for education: Foundational capabilities, potentials, and challenges](#). *ArXiv*, abs/2401.08664.
- Valentin Li’evin, Christoffer Egeberg Hother, and Ole Winther. 2022. [Can large language models reason about medical questions?](#) *Patterns*, 5.
- Fen Liu, Jianfeng Chen, Kemeng Li, Weijie Tan, Chang Cai, and Muhammad Saad Ayub. 2022. [A parallel multi-modal factorized bilinear pooling fusion method based on the semi-tensor product for emotion recognition](#). *Entropy*, 24.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. [Hierarchical question-image co-attention for visual question answering](#). *ArXiv*, abs/1606.00061.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and A. Kalyan. 2022. [Learn to explain: Multimodal reasoning via thought chains for science question answering](#). *ArXiv*, abs/2209.09513.
- Kamil Malinka, Martin Peresíni, Anton Firc, Ondřej Hujňák, and Filip Janus. 2023. [On the educational impact of chatgpt: Is artificial intelligence ready to obtain a university degree?](#) *Proceedings of the 2023 Conference on Innovation and Technology in Computer Science Education V. 1*.
- Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. 2015. [Ask your neurons: A neural-based approach to answering questions about images](#). *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1–9.
- Hyeonwoo Noh and Bohyung Han. 2016. [Training recurrent answering units with joint loss minimization for vqa](#). *ArXiv*, abs/1606.03647.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis Christiano, Jan Leike, and Ryan J. Lowe. 2022. [Training language models to follow instructions with human feedback](#). *ArXiv*, abs/2203.02155.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikanan Sankarasubbu. 2022. [Medmcqa : A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In *ACM Conference on Health, Inference, and Learning*.
- Jorge Poco and Jeffrey Heer. 2017. [Reverse-engineering visualizations: Recovering visual encodings from chart images](#). *Computer Graphics Forum*, 36.
- Ori Ram, Yuval Kirstain, Jonathan Berant, Amir Globerson, and Omer Levy. 2021. [Few-shot question answering by pretraining span selection](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Shailaja Keyur Sampat, Yezhou Yang, and Chitta Baral. 2020. [Visuo-linguistic question answering \(vlqa\) challenge](#). In *Findings*.
- Teo Susnjak. 2022. [Chatgpt: The end of online exam integrity?](#) *ArXiv*, abs/2212.09292.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. [Large language models in medicine](#). *Nature Medicine*, 29:1930–1940.
- Maheraj Thiruvanantharajah, Nawanjana Hangarangoda, and S.C. Rajapakshe. 2021. [Automated question and answer generating system for educational platforms](#). *2021 6th International*

- Conference on Information Technology Research (ICITR)*, pages 1–6.
- James Vo. 2024. [Vi-mistral-x: Building a vietnamese language model with advanced continual pre-training](#). *ArXiv*, abs/2403.15470.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Huai hsin Chi, F. Xia, Quoc Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). *ArXiv*, abs/2201.11903.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. 2023a. [Autogen: Enabling next-gen llm applications via multi-agent conversation framework](#). *ArXiv*, abs/2308.08155.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. 2023b. [Bloomberggpt: A large language model for finance](#). *ArXiv*, abs/2303.17564.
- Yiran Wu, Feiran Jia, Shaokun Zhang, Han-Tai Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2023c. [An empirical study on challenging math problem solving with gpt-4](#). *ArXiv*, abs/2306.01337.
- Qi Xia, Thomas K. F. Chiu, Xinyan Zhou, Ching Sing Chai, and Miaoting Cheng. 2022. [Systematic literature review on opportunities, challenges, and future research recommendations of artificial intelligence in education](#). *Comput. Educ. Artif. Intell.*, 4:100118.
- Ke Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In *International Conference on Machine Learning*.
- Weiwen Xu, Xin Li, Wenxuan Zhang, Meng Zhou, Lidong Bing, Wai Lam, and Luo Si. 2022. [From clozing to comprehending: Retrofitting pre-trained language model to pre-trained machine reader](#). *ArXiv*, abs/2212.04755.
- Lixiang Yan, Lele Sha, Linxuan Zhao, Yuheng Li, Roberto Martínez-Maldonado, Guanliang Chen, Xinyu Li, Yueqiao Jin, and Dragan Gaevi. 2023. [Practical and ethical challenges of large language models in education: A systematic scoping review](#). *Br. J. Educ. Technol.*, 55:90–112.
- Hongyang Yang, Xiao-Yang Liu, and Chris Wang. 2023. [Fingpt: Open-source financial large language models](#). *ArXiv*, abs/2306.06031.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2015. [Stacked attention networks for image question answering](#). *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21–29.
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, and Songfang Huang. 2023. [How well do large language models perform in arithmetic tasks?](#) *ArXiv*, abs/2304.02015.
- Fanlong Zeng, Wensheng Gan, Yongheng Wang, Ning Liu, and Philip S. Yu. 2023. [Large language models for robotics: A survey](#). *ArXiv*, abs/2311.07226.
- Wenxuan Zhang, Sharifah Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2023. [M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models](#). *ArXiv*, abs/2306.05179.