

# Exploring the Capabilities of Prompted Large Language Models in Educational and Assessment Applications

Subhankar Maity  
IIT Kharagpur, India  
subhankar.ai@kgpian.  
iitkgp.ac.in

Aniket Deroy  
IIT Kharagpur, India  
roydanik18@kgpian.  
iitkgp.ac.in

Sudeshna Sarkar  
IIT Kharagpur, India  
sudeshna@cse.iitkgp.ac.in

## ABSTRACT

In the era of generative artificial intelligence (AI), the fusion of large language models (LLMs) offers unprecedented opportunities for innovation in the field of modern education. We embark on an exploration of prompted LLMs within the context of educational and assessment applications to uncover their potential. Through a series of carefully crafted research questions, we investigate the effectiveness of prompt-based techniques in generating open-ended questions from school-level textbooks, assess their efficiency in generating open-ended questions from undergraduate-level technical textbooks, and explore the feasibility of employing a *chain-of-thought* inspired multi-stage prompting approach for language-agnostic multiple-choice question (MCQ) generation. Additionally, we evaluate the ability of prompted LLMs for language learning, exemplified through a case study in the low-resource Indian language Bengali, to explain Bengali grammatical errors. We also evaluate the potential of prompted LLMs to assess human resource (HR) spoken interview transcripts. By juxtaposing the capabilities of LLMs with those of human experts across various educational tasks and domains, our aim is to shed light on the potential and limitations of LLMs in reshaping educational practices.

## Keywords

Educational, Large language models, Prompt, Question Generation, Assessment

## 1. INTRODUCTION

In the current era of rapid technological advancement, the integration of generative AI models, particularly LLMs, represents a pivotal shift in educational practices and assessment methodologies [36, 49, 4]. These LLMs, driven by generative AI, have a profound grasp of natural language and formidable computational prowess, offering promising transformative potential in both learning facilitation and student evaluation [65]. Our study embarks on a thorough exploration of the utilization of LLMs across a spectrum of

educational and assessment contexts, with a focus on elucidating their efficacy and identifying areas ready for improvement. Our objective is to address key research questions, striving to unveil the multifaceted potential of LLMs while acknowledging the inherent complexities and challenges in their integration.

Our investigation is underscored by the use of a prompting approach, aimed at enhancing the capabilities of LLMs in subsequent tasks by providing additional information, known as a "prompt", to guide their generation process [27]. Recently, the use of prompts has gained significant attention across different natural language generation tasks, such as summarization [6], machine translation [55], etc. Through rigorous examination and analysis, our objective is to contribute meaningfully to the ongoing discourse surrounding the integration of generative AI models in education, providing nuanced insights that inform future research endeavors and educational practices.

## 2. RELATED WORK

The present research studies [18, 26] explore various prompt-based strategies for question generation (QG). [18] curate the KHANQ dataset, categorizing each data sample into a triple of <Context, Prompt, Question> and investigate prompt-based QG using LLMs such as BERT generation [47], BART [28], GPT2 [41], T5 [42], and UniLM [16]. The prompts used in KHANQ are tailored according to the learners' background knowledge and comprehension of the subject matter. Despite the considerable value of the KHANQ dataset, the authors have not made it available so far. [26] utilize prompt-based fine-tuning to formulate multi-hop questions. The methodology entails a sequence of tasks, beginning with QG and subsequently transitioning to question-answering (QA), which is iteratively performed to refine the QG process. T5 is used to train both the QG and QA models. Additionally, question paraphrasing is implemented to enhance the method's robustness. Lastly, prompt-based fine-tuning is employed to produce high-quality questions. They generated a prompt by selecting pertinent words related to the accurate answer and evaluated their model on the HotpotQA [63], SQuAD [43], and Quora Question Pairs datasets [56].

Recent studies [60, 35, 26] in automated QG that leverage LLMs have utilized single-hop QA datasets such as SQuAD and multi-hop QA datasets such as HotpotQA. These QA datasets comprise <Context, Question, Answer> triples,

S. Maity, A. Deroy, and S. Sarkar. Exploring the capabilities of prompted large language models in educational and assessment applications. In B. Paaßen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 961–968, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.

© 2024 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.  
<https://doi.org/10.5281/zenodo.12730013>

wherein *Context* denotes a contextual document, *Question* is a query formulated by a human and *Answer* is its associated response. Current QG methods have also benefited from the availability of QA datasets, such as the Natural Questions corpus [24], QuAC [12], TriviaQA [21], NewsQA [59], QG-STEC [50], etc. However, it is worth noting that we have identified several limitations in the existing datasets:

- As highlighted in [2, 11], these datasets are limited in the generation of simple factual questions.
- As indicated by [2, 31, 33], many of these QA datasets are sourced from open-access platforms, such as Wikipedia articles. Typically, they lack the educational aspect and do not require higher-order cognitive skills to answer them.

Moreover, there has been no exploration of the capabilities of prompted LLMs for generating open-ended questions from educational textbooks.

The generation of MCQs in multilingual settings, particularly for low-resource languages, is crucial to overcome language barriers, improve accessibility, and advance education in marginalized communities. Although previous research [45] has been conducted for the English language, involving fine-tuning of a T5 model on the DG-RACE dataset [25] to produce distractors for MCQ, currently there is no research available for multilingual contexts, such as German, Hindi, and Bengali, where an encoder-decoder-based model is used for distractor generation. Additionally, there is currently no research on MCQ generation that investigates the potential of the *chain-of-thought* [62] inspired prompt-based method to generate MCQs in various languages.

Despite increasing interest in grammatical error correction (GEC) and the availability of GEC datasets in major languages such as English [14, 10, 34], Chinese [67], German [8], Russian [48], Spanish [15], etc., there is a noticeable shortage of real-world GEC datasets specifically designed for low-resource languages such as Bengali (despite being the 7<sup>th</sup> most spoken language worldwide [5]). Current synthetic Bengali GEC datasets, as mentioned in [19], lack the authenticity and diversity required to represent the complexities of real-world language usage. Although there is existing GEC research for Bengali [19, 3, 20, 39, 52], no effort has been made in the domains of feedback or explanation generation within this particular context. Furthermore, there has been no investigation in GEC to assess the potential of generative pre-trained LLMs like GPT-4 Turbo, GPT-3.5 Turbo, Llama-2, etc., for low-resource languages such as Bengali.

Recent studies have explored aspects of speech scoring, such as assessing response content. This involves modeling features extracted from response transcriptions alongside the corresponding question to gauge response relevance [64, 38]. Expanding on this, [37] improved their approach by integrating acoustic cues and grammar features to enhance scoring accuracy. In a more recent investigation, [53] used speech and text transformers [51] to evaluate candidate speech. To our knowledge, no research has investigated the use of state-of-the-art LLMs for automated human resource (HR) interview evaluation. Moreover, earlier research in automated

speech scoring focused primarily on *scoring*, with minimal emphasis on *error detection* and *providing feedback along with suggestions* for improvement.

### 3. RESEARCH QUESTIONS

In this section, we present the key research questions guiding our investigation of the capabilities of prompted LLMs across diverse educational and assessment contexts. These research questions serve as focal points, with the aim of evaluating the effectiveness of LLMs compared to human experts in different tasks and domains. We address the following research questions (RQs) on a diverse set of educational topics, as described below.

- RQ1:** To what extent are prompt-based techniques [18, 26] effective in generating open-ended questions using LLMs from school-level textbooks compared to human experts?
- RQ2:** To what extent are prompt-based techniques effectively enabling LLMs to generate open-ended questions from undergraduate-level technical textbooks compared to human experts?
- RQ3:** Can a *chain-of-thought* [62] inspired multi-stage prompting approach be developed to generate language-agnostic multiple-choice questions using GPT-based models?
- RQ4:** To what extent are pre-trained LLMs capable of explaining Bengali grammatical errors compared to human experts?
- RQ5:** How ready are pre-trained LLMs to assess human resource spoken interview transcripts compared to human experts?

### 4. CURRENT RESEARCH PROGRESS

In this section, we discuss the current research progress that has been made in addressing the aforementioned research questions.

**RQ1:** *To what extent are prompt-based techniques [18, 26] effective in generating open-ended questions using LLMs from school-level textbooks compared to human experts?*

To answer this research question, we propose to examine the efficacy of prompt-based methods [18, 26] in generating open-ended questions using LLMs from school-level textbooks, compared to human experts. Prompt-based techniques entail furnishing textual cues or prompts to guide LLMs in crafting questions coherent with a given context. These prompts act as signals for the LLMs to produce relevant and coherent questions. Our study aims to investigate the effectiveness of these prompt-based techniques in generating descriptive and reasoning-based questions tailored to educational contexts.

In our methodology [29], we address the challenge posed by the inadequacy of existing QA datasets for prompt-based QG in educational settings by curating a new dataset called EduProbe. This dataset is specifically adapted for school-level subjects (e.g., history, geography, economics, environmental studies, and science) and draws on the rich content

of the NCERT<sup>1</sup> textbooks. Each instance in the dataset is annotated with quadruples comprising: 1) *Context*: a segment serving as the basis for question formulation, 2) *Long Prompt*: an extended textual cue encompassing the core theme of the context, 3) *Short Prompt*: a condensed representation of crucial information or focus within the context, and 4) *Question*: a question in line with the context and aligned with the prompts.

Different prompts not only speed up the process of creating questions but also improve the overall quality and diversity of the questions produced by providing LLMs additional guidance on what information to give more importance when creating questions. We explore various prompt-based QG techniques (e.g., *long prompt*, *short prompt*, and *without prompt*) by fine-tuning pre-trained transformer-based LLMs, including PEGASUS [66], T5 [42], and BART [28]. Furthermore, we examine the performance of two general-purpose pre-trained LLMs, text-davinci-003 [9] and GPT-3.5 Turbo, using a zero-shot prompting approach.

Through automated evaluation, we demonstrate that T5 (*with long prompt*) outperforms other LLMs, although it falls short of the human baseline. Intriguingly, text-davinci-003 consistently shows superior results compared to other LLMs in various prompt settings, even surpassing them in human evaluation criteria. However, prompt-based QG models mostly fall below the human baseline, indicating the need for further exploration and refinement in this domain.

**RQ2: To what extent are prompt-based techniques effectively enabling LLMs to generate open-ended questions from undergraduate-level technical textbooks compared to human experts?**

To address this research question, we delve into the effectiveness of prompt-based techniques in facilitating LLMs to generate open-ended questions from technical textbooks at the undergraduate level, in comparison to human experts. Our investigation focuses on the automated generation of various open-ended questions in the technical domain, an area that is relatively less explored in educational QG research [2].

To facilitate our study, we curate EngineeringQ from undergraduate level technical textbooks on subjects such as operating systems and computer networks. This dataset is designed for prompt-based QG and comprises triples consisting of 1) *Context*: segments from which questions are derived, 2) *Prompt*: concise and specific keyphrases guiding QG, and 3) *Question*: questions coherent with context and prompt.

We evaluate several fine-tuned encoder-decoder based LLMs, such as Pegasus, BART, Flan-T5 [13], and T5, on EngineeringQ. Additionally, we explore the potential of general-purpose decoder-only LLMs like GPT-3.5 Turbo, text-davinci-003, and GPT-4 [1] using a zero-shot prompting approach. Our assessment involves both automated metrics and human evaluation by domain experts. Moreover, we examine the domain adaptation [54, 2, 44] capability of LLMs by fine-

tuning the best-performing LLM on school-level subjects (e.g., history, geography, economics, environmental studies, and science) and assessing its efficacy on undergraduate-level computer science and information technology subjects (e.g., operating systems and computer networks) for zero-shot and few-shot QG. To gauge question complexity, we employ Bloom’s revised taxonomy [7], enhancing our understanding of their educational significance.

Experimental findings indicate that T5<sub>LARGE</sub> outperforms other LLMs in automated evaluation metrics, while text-davinci-003 excels in human evaluation metrics. However, LLMs in both scenarios fall short of the human baseline, highlighting the need for further refinement and exploration in this domain.

**RQ3: Can a chain-of-thought [62] inspired multi-stage prompting approach be developed to generate language-agnostic multiple-choice questions using GPT-based models?**

To answer this research question, we present a novel *chain-of-thought* inspired multi-stage prompting strategy for crafting language-agnostic MCQs utilizing GPT-based models [30]. This method, known as the multi-stage prompting approach (MSP), capitalizes on the strengths of GPT models such as text-davinci-003 and GPT-4, renowned for their proficiency across diverse natural language processing tasks. Our proposed MSP technique integrates the innovative concept of *chain-of-thought* prompting [62], wherein the GPT model receives a sequence of interconnected cues to guide the MCQ generation process.

We evaluated our proposed language-agnostic MCQ generation method on several datasets across different languages. SQuAD served as the MCQ generation dataset for English (En), while GermanQuAD [32] was utilized for German (De). For generating questions in Hindi (Hi), we employed HiQuAD [23], and for Bengali (Bn), we utilized BanglaRQA [17].

Through automated evaluation, we consistently demonstrate the superiority of the MSP method over the conventional single-stage prompting (SSP) baseline, evident in the production of high-quality distractors crucial for effective MCQs. Furthermore, our one-shot MSP method enhances automatic evaluation results, contributing to improved distractor generation in multiple languages, including English, German, Bengali, and Hindi. In human evaluation, questions generated using our proposed MSP approach exhibit superior levels of *grammaticality* [61], *answerability* [61], and *difficulty* [22] for high-resource languages (e.g., En, De), underscoring its effectiveness in diverse linguistic contexts. However, further research and fine-tuning of GPT-based models might be required to improve the results for low-resource languages (e.g., Hi, Bn) and to reduce the disparity with high-resource languages (e.g., En, De) in both automated and human evaluation criteria.

**RQ4: To what extent are pre-trained LLMs capable of explaining Bengali grammatical errors compared to human experts?**

<sup>1</sup>NCERT: <http://tinyurl.com/3x7hm2jk>

GEC tools, driven by advanced generative AI, excel at rectifying linguistic inaccuracies in user input. However, they often lack in furnishing essential natural language explanations, crucial for language learning and comprehension of grammatical rules. Particularly in low-resource Indian languages like Bengali, there is limited exploration of these tools, necessitating grammatical error explanation (GEE) systems that not only correct sentences but also provide explanations for errors.

To address this research question, we propose an investigation into the proficiency of pre-trained LLMs including GPT-4 Turbo, GPT-3.5 Turbo, text-davinci-003, text-babbage-001, text-curie-001, text-ada-001, llama-2-7b [57], llama-2-13b, and llama-2-70b in explaining Bengali grammatical errors compared to human experts.

We introduce a real-world, multi-domain dataset sourced from various domains such as Bengali essays, social media, and news, serving as an evaluation benchmark for the GEE system. This dataset facilitates the assessment of various pre-trained LLMs against human experts for performance comparison in a one-shot prompt setting.

Our methodical experimental procedure involved both LLM and human experts, performing two crucial tasks independently. First, they were tasked with producing an accurate Bengali sentence by detecting and correcting errors in the provided sentences, ensuring both grammatical correctness and contextual appropriateness. Second, for each corrected error, they were required to categorize the *error type* and offer concise *explanations* concerning the grammatical, syntactical, or semantic issues addressed.

Our research highlights the limitations in the automatic deployment of current state-of-the-art pre-trained LLMs for Bengali GEE. We advocate for human intervention, proposing the integration of manual checks to refine GEC tools in Bengali, emphasizing the educational aspect of language learning.

**RQ5: How ready are pre-trained LLMs to assess human resource spoken interview transcripts compared to human experts?**

To address this research question, we propose a detailed examination of the readiness of pre-trained LLMs in evaluating human resource (HR) spoken interview transcripts compared to human experts. Our comprehensive analysis encompasses a range of prominent pre-trained LLMs, including GPT-4 Turbo, GPT-3.5 Turbo, text-davinci-003, text-babbage-001, text-curie-001, text-ada-001, llama-2-7b, llama-2-13b, and llama-2-70b, assessing their performance in *providing scores*, *error identification*, and *offering feedback and improvement suggestions* to candidates during simulated HR interviews.

We introduce a dataset named HURIT (Human Resource Interview Transcripts), comprising HR interview transcripts collected from real-world scenarios. The dataset consists of HR interview transcripts obtained from L2 English speakers, primarily featuring interviews conducted in the Asian region. These transcripts are derived from simulated HR

interviews in which students provided their responses. The responses were captured in *.mp3* format and subsequently transcribed into text using OpenAI’s Whisper large-v2 model [40]. This dataset facilitates the evaluation of various pre-trained LLMs against human experts for performance comparison in a zero-shot prompt setting.

Our approach involved a structured assessment procedure in which both LLMs and human assessors independently *scored*, *identified errors*, and *provided constructive feedback* on HR interview transcripts. This comprehensive method enabled a thorough evaluation of each LLM’s performance, encompassing its *scoring accuracy*, *error detection*, and *feedback provision*. Additionally, we compared their abilities with those of expert human evaluators on various human evaluation criteria, such as *fluency*, *coherence*, *tone/politeness*, *relevance*, *conciseness*, and *grammaticality* [58, 46].

Our findings highlight the proficiency of pre-trained LLMs, particularly GPT-4 Turbo and GPT-3.5 Turbo, in delivering evaluations comparable to those provided by expert human evaluators. However, while these LLMs excel in *scoring* candidates, they often struggle to *identify errors* and *provide actionable feedback* for performance improvement in HR interviews. Our research underscores that although pre-trained LLMs demonstrate promise in certain aspects, they are not yet fully equipped for automatic deployment in HR interview assessments. Instead, we advocate for a human-in-the-loop approach, emphasizing the importance of manual checks to address inconsistencies and improve the quality of feedback provided, presenting a more viable strategy for HR interview assessment.

## 5. CONCLUSION

Our study addressed key research questions about the integration of LLMs in educational and assessment applications. We investigated the effectiveness of prompt-based techniques in generating open-ended questions from school-level textbooks using LLMs, highlighting promising but imperfect performance compared to human experts. Despite advancements, LLMs struggled to match human expertise in generating open-ended questions from undergraduate-level technical textbooks, indicating areas for improvement. Additionally, our proposed MSP approach for crafting language-agnostic MCQs shows that further research and fine-tuning of GPT models are required to improve the results in low-resource languages (e.g., Hi, Bn). Furthermore, our exploration of the ability of LLMs to *explain* Bengali grammatical errors revealed deficiencies, underscoring the importance of human intervention. Lastly, while LLMs showed competence in *scoring* HR interview transcripts, they encountered challenges in *error identification* and *feedback provision*, emphasizing the need for human oversight. Overall, our study underscores the potential of LLMs in educational and assessment applications but highlights the ongoing need for research and refinement to fully harness their capabilities.

In the doctoral consortium, we anticipate receiving recommendations and feedback regarding the current status of our research progress.

## 6. REFERENCES

- [1] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] S. Al Faraby, A. Adiwijaya, and A. Romadhony. Review on neural question generation for education purposes. *International Journal of Artificial Intelligence in Education*, pages 1–38, 2023.
- [3] P. Bagchi, M. Arafin, A. Akther, and K. M. Alam. Bangla spelling error detection and correction using n-gram model. In *International Conference on Machine Intelligence and Emerging Technologies*, pages 468–482. Springer, 2022.
- [4] Z. Bahroun, C. Anane, V. Ahmed, and A. Zacca. Transforming education: A comprehensive review of generative artificial intelligence in educational settings through bibliometric and content analysis. *Sustainability*, 15(17):12983, 2023.
- [5] E. Behrman, A. Santra, S. Sarkar, P. Roy, R. Yadav, S. Dutta, and A. Ghosal. Dialect identification of the bengali language. In *Data Science and Data Analytics*, pages 357–373. Chapman and Hall/CRC, 2021.
- [6] A. Bhaskar, A. Fabbri, and G. Durrett. Prompted opinion summarization with GPT-3.5. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9282–9300, Toronto, Canada, July 2023. Association for Computational Linguistics.
- [7] B. S. Bloom. *A taxonomy for learning, teaching, and assessing: A revision of Bloom’s taxonomy of educational objectives*. Longman, 2010.
- [8] A. Boyd. Using Wikipedia edits in low resource grammatical error correction. In W. Xu, A. Ritter, T. Baldwin, and A. Rahimi, editors, *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-generated Text*, pages 79–84, Brussels, Belgium, Nov. 2018. Association for Computational Linguistics.
- [9] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [10] C. Bryant, M. Felice, Ø. E. Andersen, and T. Briscoe. The BEA-2019 shared task on grammatical error correction. In H. Yannakoudakis, E. Kochmar, C. Leacock, N. Madnani, I. Pilán, and T. Zesch, editors, *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75, Florence, Italy, Aug. 2019. Association for Computational Linguistics.
- [11] S. Cao and L. Wang. Controllable open-ended question generation with a new question type ontology. In C. Zong, F. Xia, W. Li, and R. Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6424–6439, Online, Aug. 2021. Association for Computational Linguistics.
- [12] E. Choi, H. He, M. Iyyer, M. Yatskar, W.-t. Yih, Y. Choi, P. Liang, and L. Zettlemoyer. QuAC: Question answering in context. In E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium, Oct.-Nov. 2018. Association for Computational Linguistics.
- [13] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, Y. Li, X. Wang, M. Dehghani, S. Brahma, A. Webson, S. S. Gu, Z. Dai, M. Suzgun, X. Chen, A. Chowdhery, A. Castro-Ros, M. Pellat, K. Robinson, D. Valter, S. Narang, G. Mishra, A. Yu, V. Zhao, Y. Huang, A. Dai, H. Yu, S. Petrov, E. H. Chi, J. Dean, J. Devlin, A. Roberts, D. Zhou, Q. V. Le, and J. Wei. Scaling instruction-finetuned language models, 2022.
- [14] D. Dahlmeier, H. T. Ng, and S. M. Wu. Building a large annotated corpus of learner English: The NUS corpus of learner English. In J. Tetreault, J. Burstein, and C. Leacock, editors, *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 22–31, Atlanta, Georgia, June 2013. Association for Computational Linguistics.
- [15] S. Davidson, A. Yamada, P. Fernandez Mira, A. Carando, C. H. Sanchez Gutierrez, and K. Sagae. Developing NLP tools with a new corpus of learner Spanish. In N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, and S. Piperidis, editors, *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 7238–7243, Marseille, France, May 2020. European Language Resources Association.
- [16] L. Dong, N. Yang, W. Wang, F. Wei, X. Liu, Y. Wang, J. Gao, M. Zhou, and H.-W. Hon. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 32, 2019.
- [17] S. M. S. Ekram, A. A. Rahman, M. S. Altaf, M. S. Islam, M. M. Rahman, M. M. Rahman, M. A. Hossain, and A. R. M. Kamal. BanglaRQA: A benchmark dataset for under-resourced Bangla language reading comprehension-based question answering with diverse question-answer types. In Y. Goldberg, Z. Kozareva, and Y. Zhang, editors, *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2518–2532, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics.
- [18] H. Gong, L. Pan, and H. Hu. KHANQ: A dataset for generating deep questions in education. In N. Calzolari, C.-R. Huang, H. Kim, J. Pustejovsky, L. Wanner, K.-S. Choi, P.-M. Ryu, H.-H. Chen, L. Donatelli, H. Ji, S. Kurohashi, P. Paggio, N. Xue, S. Kim, Y. Hahm, Z. He, T. K. Lee, E. Santus,

- F. Bond, and S.-H. Na, editors, *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5925–5938, Gyeongju, Republic of Korea, Oct. 2022. International Committee on Computational Linguistics.
- [19] N. Hossain, M. H. Bijoy, S. Islam, and S. Shatabda. Panini: a transformer-based grammatical error correction method for bangla. *Neural Computing and Applications*, pages 1–15, 2023.
- [20] N. Hossain, S. Islam, and M. N. Huda. Development of bangla spell and grammar checkers: Resource creation and evaluation. *IEEE Access*, 9:141079–141097, 2021.
- [21] M. Joshi, E. Choi, D. Weld, and L. Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In R. Barzilay and M.-Y. Kan, editors, *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada, July 2017. Association for Computational Linguistics.
- [22] D. Kalpakchi and J. Boye. Quasi: a synthetic question-answering dataset in Swedish using GPT-3 and zero-shot learning. In T. Alumäe and M. Fishel, editors, *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 477–491, Tórshavn, Faroe Islands, May 2023. University of Tartu Library.
- [23] V. Kumar, N. Joshi, A. Mukherjee, G. Ramakrishnan, and P. Jyothi. Cross-lingual training for automatic question generation. In A. Korhonen, D. Traum, and L. Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4863–4872, Florence, Italy, July 2019. Association for Computational Linguistics.
- [24] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, J. Devlin, K. Lee, K. Toutanova, L. Jones, M. Kelcey, M.-W. Chang, A. M. Dai, J. Uszkoreit, Q. Le, and S. Petrov. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466, 2019.
- [25] G. Lai, Q. Xie, H. Liu, Y. Yang, and E. Hovy. RACE: Large-scale Reading comprehension dataset from examinations. In M. Palmer, R. Hwa, and S. Riedel, editors, *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark, Sept. 2017. Association for Computational Linguistics.
- [26] S. Lee and M. Lee. Type-dependent prompt CycleQAG : Cycle consistency for multi-hop question generation. In N. Calzolari, C.-R. Huang, H. Kim, J. Pustejovsky, L. Wanner, K.-S. Choi, P.-M. Ryu, H.-H. Chen, L. Donatelli, H. Ji, S. Kurohashi, P. Paggio, N. Xue, S. Kim, Y. Hahm, Z. He, T. K. Lee, E. Santus, F. Bond, and S.-H. Na, editors, *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6301–6314, Gyeongju, Republic of Korea, Oct. 2022. International Committee on Computational Linguistics.
- [27] B. Lester, R. Al-Rfou, and N. Constant. The power of scale for parameter-efficient prompt tuning. In M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics.
- [28] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In D. Jurafsky, J. Chai, N. Schuster, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July 2020. Association for Computational Linguistics.
- [29] S. Maity, A. Deroy, and S. Sarkar. Harnessing the power of prompt-based techniques for generating school-level questions using large language models. In *Proceedings of the 15th Annual Meeting of the Forum for Information Retrieval Evaluation, FIRE '23*, page 30–39, New York, NY, USA, 2024. Association for Computing Machinery.
- [30] S. Maity, A. Deroy, and S. Sarkar. A novel multi-stage prompting approach for language agnostic mcq generation using gpt. In N. Goharian, N. Tonello, Y. He, A. Lipani, G. McDonald, C. Macdonald, and I. Ounis, editors, *Advances in Information Retrieval*, pages 268–277, Cham, 2024. Springer Nature Switzerland.
- [31] R. Mitkov, H. Maslak, T. Ranasinghe, V. Sosoni, et al. Automatic generation of multiple-choice test items from paragraphs using deep neural networks. In *Advancing Natural Language Processing in Educational Assessment*, pages 77–89. Routledge, 2023.
- [32] T. Möller, J. Risch, and M. Pietsch. GermanQuAD and GermanDPR: Improving non-English question answering and passage retrieval. In A. Fisch, A. Talmor, D. Chen, E. Choi, M. Seo, P. Lewis, R. Jia, and S. Min, editors, *Proceedings of the 3rd Workshop on Machine Reading for Question Answering*, pages 42–50, Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics.
- [33] N. Mulla and P. Gharpure. Automatic question generation: a review of methodologies, datasets, evaluation metrics, and applications. *Progress in Artificial Intelligence*, 12(1):1–32, 2023.
- [34] C. Napoles, K. Sakaguchi, and J. Tetreault. JFLEG: A fluency corpus and benchmark for grammatical error correction. In M. Lapata, P. Blunsom, and A. Koller, editors, *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 229–234, Valencia, Spain, Apr. 2017. Association for Computational Linguistics.
- [35] L. Pan, Y. Xie, Y. Feng, T.-S. Chua, and M.-Y. Kan. Semantic graphs for generating deep questions. In D. Jurafsky, J. Chai, N. Schuster, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1463–1475, Online, July 2020. Association for

Computational Linguistics.

- [36] J. Prather, P. Denny, J. Leinonen, B. A. Becker, I. Albluwi, M. Craig, H. Keuning, N. Kiesler, T. Kohn, A. Luxton-Reilly, S. MacNeil, A. Petersen, R. Pettit, B. N. Reeves, and J. Savelka. The robots are here: Navigating the generative ai revolution in computing education. In *Proceedings of the 2023 Working Group Reports on Innovation and Technology in Computer Science Education, ITiCSE-WGR '23*, page 108–159, New York, NY, USA, 2023. Association for Computing Machinery.
- [37] Y. Qian, P. Lange, K. Evanini, R. Pugh, R. Ubale, M. Mulholland, and X. Wang. Neural approaches to automated speech scoring of monologue and dialogue responses. In *ICASSP 2019-2019 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 8112–8116. IEEE, 2019.
- [38] Y. Qian, R. Ubale, M. Mulholland, K. Evanini, and X. Wang. A prompt-aware neural network approach to content-based scoring of non-native spontaneous speech. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 979–986, 2018.
- [39] R. Z. Rabbi, M. I. R. Shuvo, and K. A. Hasan. Bangla grammar pattern recognition using shift reduce parser. In *2016 5th International Conference on Informatics, Electronics and Vision (ICIEV)*, pages 229–234, 2016.
- [40] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever. Robust speech recognition via large-scale weak supervision, 2022.
- [41] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [42] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [43] P. Rajpurkar, R. Jia, and P. Liang. Know what you don't know: Unanswerable questions for SQuAD. In I. Gurevych and Y. Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia, July 2018. Association for Computational Linguistics.
- [44] M. Reid, V. Zhong, S. Gururangan, and L. Zettlemoyer. M2D2: A massively multi-domain language modeling dataset. In Y. Goldberg, Z. Kozareva, and Y. Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 964–975, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics.
- [45] R. Rodriguez-Torrealba, E. Garcia-Lopez, and A. Garcia-Cabot. End-to-end generation of multiple-choice questions using text-to-text transfer transformer models. *Expert Systems with Applications*, 208:118258, 2022.
- [46] S. J. Ross. *Interviewing for language proficiency*. Springer, 2017.
- [47] S. Rothe, S. Narayan, and A. Severyn. Leveraging pre-trained checkpoints for sequence generation tasks. *Transactions of the Association for Computational Linguistics*, 8:264–280, 2020.
- [48] A. Rozovskaya and D. Roth. Grammar error correction in morphologically rich languages: The case of Russian. *Transactions of the Association for Computational Linguistics*, 7:1–17, 2019.
- [49] L. I. Ruiz-Rojas, P. Acosta-Vargas, J. De-Moreta-Llovet, and M. Gonzalez-Rodriguez. Empowering education with generative artificial intelligence tools: Approach with an instructional design matrix. *Sustainability*, 15(15), 2023.
- [50] V. Rus, B. Wyse, P. Piwek, M. Lintean, S. Stoyanchev, and C. Moldovan. The first question generation shared task evaluation challenge. In J. Kelleher, B. M. Namee, and I. v. d. Sluis, editors, *Proceedings of the 6th International Natural Language Generation Conference*. Association for Computational Linguistics, July 2010.
- [51] J. Shah, Y. K. Singla, C. Chen, and R. R. Shah. What all do audio transformer models hear? probing acoustic representations for language delivery and its structure. *arXiv preprint arXiv:2101.00387*, 2021.
- [52] S. F. Shetu, M. Saifuzzaman, M. Parvin, N. N. Moon, R. Yousuf, and S. Sultana. Identifying the writing style of bangla language using natural language processing. In *2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–6, 2020.
- [53] Y. K. Singla, A. Gupta, S. Bagga, C. Chen, B. Krishnamurthy, and R. R. Shah. Speaker-conditioned hierarchical modeling for automated speech scoring. In *Proceedings of the 30th ACM international conference on information & knowledge management*, pages 1681–1691, 2021.
- [54] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu. A survey on deep transfer learning. In *Artificial Neural Networks and Machine Learning–ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part III 27*, pages 270–279. Springer, 2018.
- [55] Z. Tan, X. Zhang, S. Wang, and Y. Liu. MSP: Multi-stage prompting for making pre-trained language models better translators. In S. Muresan, P. Nakov, and A. Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6131–6142, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [56] N. Thakur, N. Reimers, J. Daxenberger, and I. Gurevych. Augmented SBERT: Data augmentation method for improving bi-encoders for pairwise sentence scoring tasks. In K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, R. Cotterell, T. Chakraborty, and Y. Zhou, editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 296–310, Online, June 2021. Association for Computational Linguistics.
- [57] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C.

- Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
- [58] N. Travers and L.-S. Huang. Breaking Intangible Barriers in English-as-an-Additional-Language Job Interviews: Evidence from Interview Training and Ratings. *Applied Linguistics*, 42(4):641–667, 11 2020.
- [59] A. Trischler, T. Wang, X. Yuan, J. Harris, A. Sordoni, P. Bachman, and K. Suleman. NewsQA: A machine comprehension dataset. In P. Blunsom, A. Bordes, K. Cho, S. Cohen, C. Dyer, E. Grefenstette, K. M. Hermann, L. Rimell, J. Weston, and S. Yih, editors, *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200, Vancouver, Canada, Aug. 2017. Association for Computational Linguistics.
- [60] L. A. Tuan, D. Shah, and R. Barzilay. Capturing greater context for question generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9065–9072, Apr. 2020.
- [61] A. Ushio, F. Alva-Manchego, and J. Camacho-Collados. Generative language models for paragraph-level question generation. In Y. Goldberg, Z. Kozareva, and Y. Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 670–688, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics.
- [62] J. Wei, X. Wang, D. Schuurmans, M. Bosma, b. ichter, F. Xia, E. Chi, Q. V. Le, and D. Zhou. Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc., 2022.
- [63] Z. Yang, P. Qi, S. Zhang, Y. Bengio, W. Cohen, R. Salakhutdinov, and C. D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium, Oct.–Nov. 2018. Association for Computational Linguistics.
- [64] S.-Y. Yoon and C. M. Lee. Content modeling for automated oral proficiency scoring system. In H. Yannakoudakis, E. Kochmar, C. Leacock, N. Madnani, I. Pilán, and T. Zesch, editors, *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 394–401, Florence, Italy, Aug. 2019. Association for Computational Linguistics.
- [65] H. Yu. The application and challenges of chatgpt in educational transformation: New demands for teachers’ roles. *Heliyon*, 10(2):e24289, 2024.
- [66] J. Zhang, Y. Zhao, M. Saleh, and P. Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR, 2020.
- [67] Y. Zhang, Z. Li, Z. Bao, J. Li, B. Zhang, C. Li, F. Huang, and M. Zhang. MuCGEC: a multi-reference multi-source evaluation dataset for Chinese grammatical error correction. In M. Carpuat, M.-C. de Marneffe, and I. V. Meza Ruiz, editors, *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3118–3130, Seattle, United States, July 2022. Association for Computational Linguistics.