

How Teachers Can Use Large Language Models and Bloom’s Taxonomy to Create Educational Quizzes

Sabina Elkins^{1,2}, Ekaterina Kochmar^{2,3}, Jackie C.K. Cheung^{1,4}, Iulian Serban²

¹McGill University & MILA

²Korbit Technologies Inc.

³MBZUAI

⁴Canada CIFAR AI Chair

Abstract

Question generation (QG) is a natural language processing task with an abundance of potential benefits and use cases in the educational domain. In order for this potential to be realized, QG systems must be designed and validated with pedagogical needs in mind. However, little research has assessed or designed QG approaches with the input from real teachers or students. This paper applies a large language model-based QG approach where questions are generated with learning goals derived from Bloom’s taxonomy. The automatically generated questions are used in multiple experiments designed to assess how teachers use them in practice. The results demonstrate that teachers prefer to write quizzes with automatically generated questions, and that such quizzes have no loss in quality compared to handwritten versions. Further, several metrics indicate that automatically generated questions can even improve the quality of the quizzes created, showing the promise for large scale use of QG in the classroom setting.

Introduction

Question generation (QG) is a popular natural language processing (NLP) task. The goal is to generate natural-language questions that are useful and fluent. Many approaches also attempt to generate the corresponding answers, or use the answer to generate the question (Kurdi et al. 2020; Mulla and Gharpure 2023). Due to their recent success in NLP, recent QG research has been dominated by the use of Transformer-based large language models (LLMs) (Kurdi et al. 2020; Liu et al. 2023).

An obvious use case for QG is educational applications. A robust QG system could, for example, reduce the time spent by teachers to create educational content such as homework, quizzes, tests, in-class learning activities, and more. Alternatively, it could serve as a practice tool for students. The range of potential uses for educational question generation (EQG) is expansive, especially considering the recent success of LLMs (Kasneci et al. 2023; Kurdi et al. 2020).

Unfortunately, there is minimal documented real-world deployment of such systems (Kasneci et al. 2023; Kurdi et al. 2020). Potential reasons for this lack of adoption may include poor performance of older approaches, rigidity of the systems, and mistrust of users. Wang et al. (2022a) conduct a need-finding study with the aim to explore why QG systems are not being used in classrooms. One of their key findings is that QG systems must meet the needs of the educators who are using them in order to be effective and adopted. To achieve this, it is imperative that research in developing educational QG systems takes into account the opinions of their end users.

Generate questions in each level of Bloom’s taxonomy.

```

Passage: {example_context}
Remembering = {example_question}
Understanding = {example_question}
Applying = {example_question}
Analyzing = {example_question}
Evaluating = {example_question}
Creating = {example_question}

```

```

Passage: {context}
Remembering =

```

(a) *Controlled* prompt template. Due to space limits, a one-shot template is demonstrated here. In actuality, five examples (with different contexts and questions) are used.

Candidates generated using a context about convergent evolution:
Remembering: What is convergent evolution?

Applying: Can you provide an example of convergent evolution?

Analyzing: What is the difference between analogous and homologous structures or traits?

(b) *Controlled* prompting strategy generation examples.¹ All 6 levels of Bloom’s taxonomy have generations for each input context, omitted here to save space.

Figure 1: *Controlled* prompting strategy.

Given the success of LLMs in other tasks, our hypothesis is that they can generate different types of questions from a given context that teachers find useful for creating a quiz with quality comparable to a handwritten version. Further, we predict that teachers will find the generated candidates more useful when they are generated to correspond with the levels of Bloom’s taxonomy (Krathwohl 2002). Figure 1 depicts our few-shot prompting strategy to generate educational questions corresponding to Bloom’s taxonomy (more details about this approach can be found in Section). In order to evaluate our predictions, we conducted quiz writing experiments designed to compare and contrast three different quiz types. Multiple aspects of ‘usefulness’ of a quiz and quiz writing approach are considered, including *the quality of the resulting quiz*, *the efficiency* (w.r.t. time), and *the teacher’s preferences*. Our results show that the three types of quizzes are of similar quality. Some metrics demonstrate an improvement in quality when automatically generated questions are used. We also find that teachers have a strong preference for writing quizzes with the help of generated questions corresponding to Bloom’s taxonomy. These results demonstrate the huge potential of EQG for real-world classrooms and the importance of considering the needs of teachers when designing an EQG approach.

Background

Recent research in EQG, and QG more generally, revolves around the use of Transformer-based LLMs. These LLMs are deep learning models trained on massive corpora of data to improve their generative performance (Zhang et al. 2022). The reason for applying this approach in QG research is in large part due to its significant performance improvements over earlier rule-based and other types of systems (Kurdi et al. 2020; Steuer et al. 2021; Mulla and Gharpure 2023).

The typical training goal for Transformer-based LLMs is next-token prediction, meaning that they learn to predict a probable completion to an initial input text. Recent models have begun to also include reinforcement learning in their training procedure. This is the case for GPT-3.5, which is the LLM used in the experiments reported in this paper. Fine-tuning with reinforcement learning from human feedback allows GPT-3.5 to outperform its predecessors in the GPT family (Ouyang et al. 2022).

Aligning with the common LLM training objective of the next-token prediction, the emerging paradigm for QG is to feed a textual input, called a *prompt*, to an LLM for the model to complete (Mulla and Gharpure 2023). Designing this prompt to generate a desired output can be a difficult task, which has resulted in a new research direction called *prompt engineering*. One of the most common approaches to prompt engineering involves prepending a string to the context given to an LLM for generation, which is called a *prefix style prompt* (Liu et al. 2023). For instance, say a machine learning teacher wished to generate questions about gradient descent. A simple strategy they could apply is to prompt an LLM with the input:

Generate a question about gradient descent.

To increase the specificity of the generated questions, the

teacher could provide more context. For example, they might instead write a prompt with a textbook passage about a specific aspect of gradient descent, such as:

Generate a question from the following passage:
<...>

To further control generation, the teacher’s input could contain a *control element* – a keyword that will guide the generation (Mulla and Gharpure 2023). For instance, they could prompt an LLM with:

Generate a multiple-choice question from the following passage: <...>

In this simple example, we have presented three different prompts, all of which can produce different questions. Adding in different word choices and other strategies for controlling the generation quickly makes prompt engineering a complex problem to optimize.

Another aspect in prompt engineering is the inclusion of examples of the desired output format and style within the prompt itself. This is often called *few-shot learning*. In brief, few-shot learning prompts consist of an instruction, a few examples, and the task at hand. The examples are used to adapt LLMs to unseen scenarios without additional training or fine-tuning (Liu et al. 2023). For instance, following the earlier example, a teacher might prompt an LLM to generate a true-or-false type of question by including examples of such questions.

The ability to add more specificity to educational QG has allowed researchers to generate questions at different levels of difficulty, with different pedagogical goals in mind, and more. For example, Wang et al. (2022b) try a collection of different prompting strategies in an effort to optimize educational QG. They conclude that using shorter input contexts and few-shot learning results in higher quality candidate questions. Recently, Elkins et al. (2023) have demonstrated how to generate questions at different levels of question taxonomies, which are organizational structures taken from pedagogical literature, such as Bloom’s taxonomy of learning goals. The authors demonstrate successful generation of questions at various complexity levels and with different learning goals.

Despite these exciting recent research developments with educational QG, there are only a few documented cases of these techniques being used in real-world classrooms (Kasneci et al. 2023; Kurdi et al. 2020). As previously mentioned, Wang et al. (2022a) find that the lack of alignment between research goals and what teachers actually need and want from EQG is the reason that such systems are not deployed in real-world classrooms. Prior work aiming to explore the needs, opinions, and attitude of teachers and students towards using automatically generated content is few and far between. Nevertheless, there exist a few relevant precedent papers:

- The work by Van Campenhout, Hubertz, and Johnson (2022) outlines a NLP system for translating textbooks into interactive courseware. The authors conduct a large user study which shows equal student performance on machine-generated and human-written questions. However, their QG system is mostly rule-based, and their

generated questions are relatively simple (i.e., concept-matching and fill-in-the-blank).

- The aforementioned work by Elkins et al. (2023) evaluated their generated candidate questions with real teachers in an effort to more accurately assess their pedagogical usefulness. They find that their generated questions are highly rated by teachers. However, this work only evaluates questions at the individual question-level rather than at a quiz-level.
- The work by Laban et al. (2022) moves beyond question-level evaluations to a quiz writing task, similar to our work in this paper. The authors design a task where teachers make a quiz exclusively with candidate questions that were automatically generated. The teachers also marked the candidates as acceptable or not as they went; however, the final global acceptance rate was only 52%. Thus, while the authors take important steps towards the evaluation of QG in a realistic scenario, their generations themselves appear to leave room for improvement.

Methods

The EQG in this work was conducted by prompting GPT-3.5 to generate questions from a given input passage. Two different strategies are used in order to compare the pedagogically designed generation approach to a more simplistic approach to QG. The strategies are referred to as *controlled* and *simple*, respectively. Individually, both types of candidate generations are supplied to teachers in controlled experiments to assess their usefulness in practice (the details of these experiments are explained in Section). The following subsections will cover details of the input contexts, and the two prefix style prompting strategies used.

Contexts

Candidate educational questions were generated from a set of 24 passages sourced from Wikipedia, with each passage containing 5 context paragraphs, 6 to 9 sentences in length. The total number of input contexts used to generate questions is, thus, 120. The length of these contexts was determined by empirical results in preliminary work. This set of passages was manually gathered. The specific Wikipedia articles were generally chosen via hyperlinks from the domain’s main Wikipedia or glossary page in order to ensure they were relevant to the foundations of the given domain. Two domains were used: biology (BIO) and machine learning (ML). More than one domain is included in order to take steps towards demonstrating domain-agnostic results. Each domain has 12 passages, totaling 60 input contexts. The contexts underwent minor pre-processing before use: this included removal of citations, hyperlinks, footnotes and phonetic spellings, re-formatting of full sentence bullet-point lists into paragraphs, and other minor data cleaning steps.¹

¹ The input contexts, more details about Bloom’s taxonomy, the human authored few-shot examples, all of the generated candidates and quizzes, the annotator demographics, and more can be found at https://anonymous.4open.science/r/EQG_in_practice-2752/README.md.

Simple Prompting Strategy

The *simple* prompting strategy uses a generic strategy to generate questions with GPT-3.5² in an attempt to assess how well the model can generate pedagogically useful questions without any additional prompt engineering. The prompt template can be seen in Figure 2. The number of questions generated at once is six in order to produce the same amount of questions as the *controlled* strategy described below. Empirical results from preliminary experimentation showed that generating all of the questions together produced more diverse outputs, whereas generating them separately produced duplicate questions.

```
Generate 6 questions.
```

```
Passage: {context}
```

```
Questions:
```

(a) *Simple* prompt template.

```
Candidates generated using a context about
convergent evolution:
```

```
What is the definition of convergent
evolution?
```

```
What are some common functions found in
bird, bat, and pterosaur wings?
```

(b) *Simple* prompting strategy generation examples.¹ There are actually 6 generations, omitted here to save space.

Figure 2: *Simple* prompting strategy.

Controlled Prompting Strategy

The *controlled* prompting strategy uses a pedagogical question taxonomy to generate questions with different learning goals in mind. Bloom’s taxonomy is a popular framework for categorizing learning objectives in educational materials (Krathwohl 2002). The taxonomy contains six levels of learning, arranged in a hierarchical order from ‘lower’- to ‘higher’-level thinking skills: *remembering*, *understanding*, *applying*, *analyzing*, *evaluating*, and *creating*.¹ Bloom’s taxonomy helps teachers design instructional content that targets specific learning goals.

The generation strategy used in this paper is designed to reduce the overlap between generated candidates. Rather than generating one question and taxonomic level at a time, all six questions for a given context are generated at once, as shown in Figure 1. In preliminary experimentation, this approach empirically showed a reduction in the generation of identical questions and a greater diversity and adherence to the taxonomic levels.

The *controlled* prompting strategy also uses few-shot learning. Following success with preliminary experimentation, five-shot learning (i.e., five examples within the prompt) is used. As seen in Figure 1, a single example contains a prompt and six questions, one for each level of

²The model `text-davinci-003` was accessed through the OpenAI API. The dates on which the model was queried are available in the supplementary material.

Bloom’s taxonomy. The examples were handcrafted by a domain expert, and reviewed by another expert for both question quality and adherence to the intended taxonomic level. Each domain had five sets of contexts and examples, totaling 10 contexts and 60 questions.¹

Evaluation

In order to assess the applicability of our two varieties of automatically generated questions to a real-world academic setting, we must actually ask real teachers to use them. Therefore, we designed an experiment to replicate a teacher’s creation of a reading quiz. We wanted to compare the resulting quizzes’ quality when teachers are writing quizzes from scratch and when they have access to automatically generated questions while writing the quiz. Section introduces the measures of quiz quality considered in this paper. Then, Section explains the quiz writing experiments conducted by real teachers.

Quiz Quality

To be able to compare the quality of questions and quizzes written with the aid of GPT-3.5 to handwritten ones, we must first establish how to measure the quality of a quiz in the first place. However, the quality of a quiz is often a subjective metric. Teachers have differing opinions on what makes a good quiz based on their individual teaching styles. We defined a set of metrics to measure individual aspects of a quiz that capture different aspects of its quality, with only *usefulness* designed to purely reflect the annotator’s opinion.

Based on previous research, we identified the following four quiz-level metrics to include.³ A good reading quiz will be relevant to the teaching material. It will be natural and coherent, meaning that it will not confuse a student taking it. Above all, a good quiz will be approved by a teacher for classroom use. The following metrics attempt to assess each of these aspects:

- *Coverage* is a numerical metric in $[0, 1]$ which measures how much of the input passage is reflected in the final quiz. To measure this, we mapped each question to the sentences in the passage that contained pertinent information to the answer. For consistency, all questions are mapped to any possible answer in the text (i.e., both instances of repeated information are selected, open-ended questions select more of the text, etc.). A *coverage* ratio is then calculated from the length of the mapped text and the whole passage’s length. This metric is inspired by the pyramid method for annotating summaries, which uses a similar strategy (Nenkova and Passonneau 2004).
- *Structure* is an ordinal metric from 1 to 3 which measures whether the set of questions make sense together. In other words, if they are intuitively linked together with a natural/understandable flow (e.g., from easy to difficult, or from start to finish of the context). Previously, a similar metric has been used for conversational QG where questions must be logically linked in order for a conversation to be natural (Mulla and Gharpure 2023).

³The ordinal metrics also have definitions for each category available in the supplementary material.

- *Redundancy* is an ordinal metric from 1 to 3 which measures if there is redundancy/repetition within the quiz, for example if there are two questions that ask for the same answer without any different perspective or thought process required from the student. Previously, a similar metric has been used for conversational QG where questions must not be repetitive in order for a conversation to be natural (Mulla and Gharpure 2023).
- *Usefulness* is an ordinal metric from 1 to 4 which measures if a teacher would use the quiz in an assessment they create for their own class. Note that the quiz does not necessarily need to be entirely answerable from the context in order to be considered useful. Previously, similar ordinal metrics have been used in Elkins et al. (2023), and on different scales in Steuer et al. (2021) and Mulla and Gharpure (2023).

Beyond the quality of the whole quiz, it is important to ensure that the individual questions are also of high quality. Thus we also outlined three question-level metrics, meaning they are evaluated for each question in a quiz. This is not an exhaustive evaluation, but the following covers basic aspects of a question’s quality:

- *Relevancy* is a binary metric which measures whether the question is semantically relevant to the input context. Previously, similar binary approaches have been used in Steuer et al. (2021) and Elkins et al. (2023), and on different scales by Mulla and Gharpure (2023).
- *Fluency* is a binary metric which measures whether the set of questions are grammatically correct and use clear language. Similarly to the previous metric, previous approaches by Mazidi and Nielsen (2014) and Elkins et al. (2023) have applied this binary metric, and it has been used more generally on different scales in Mulla and Gharpure (2023).
- *Answerability* is a binary metric which measures whether the question can be answered from the input context. It is not necessary to be able to find a passage from the input that is an answer to the question; it is enough if a student could reasonably answer the question from the context (for example, applying logic explained in the passage to a new situation makes the question ‘answerable’). As above, previous work by Steuer et al. (2021) and Elkins et al. (2023) uses a similar binary metric, and Mulla and Gharpure (2023) suggests similar metrics on different scales.

Quiz Writing Experiments

The quiz writing experiments were designed to mimic a teacher’s creation of a reading quiz. In order to be able to measure and compare teacher’s quiz writing processes, the setting was tightly controlled. Discussion of the potential limitations this may introduce can be found in Section . Before conducting the experiments, a pilot with four teachers in the ML domain was conducted to ensure the metrics and annotator training were unambiguous. This pilot resulted in minor changes to the wording in the training and metric definitions, but no major experimental flaws were discovered.

There were 24 quiz writing teachers, 12 in each domain. The BIO teachers were found through the freelance platform Upwork and have at least a high-school level of teaching/tutoring experience. The ML teachers were recruited through word-of-mouth at the institutions of the first author and have at least a university level of teaching/tutoring experience. All teachers were fairly monetarily compensated and signed a consent form before participating. They are all proficient in the English language, and are from relatively diverse demographics.¹ The teachers completed a training module where they were provided with example quizzes in their domain and were guided through the creation of each quiz type to ensure they understood the task at hand.

The process a single teacher underwent, irrespective of their domain, is depicted in Figure 3. Each teacher wrote three quizzes, between five and ten questions in length. With 12 teachers and 12 passages per domain, each passage was used to create each type of quiz. Note that each teacher received three different passages, to reduce the potential bias due to a teacher working with material they have already seen. The three quiz types were: *handwritten*, *simple*, and *controlled*. To create a *handwritten quiz*, the teacher simply read the passage and wrote a quiz from scratch. To create a *simple quiz*, the teacher read the passage and related candidates generated with the *simple* prompting strategy. They then created a quiz with the freedom to copy generated questions directly, copy and alter them, or write questions from scratch. Similarly, to create a *controlled quiz*, the teacher read the passage and related *controlled* candidates, and then wrote a quiz. The quizzes were written in a random order to reduce any potential for biased results from the ordering of the quiz writing subtasks.

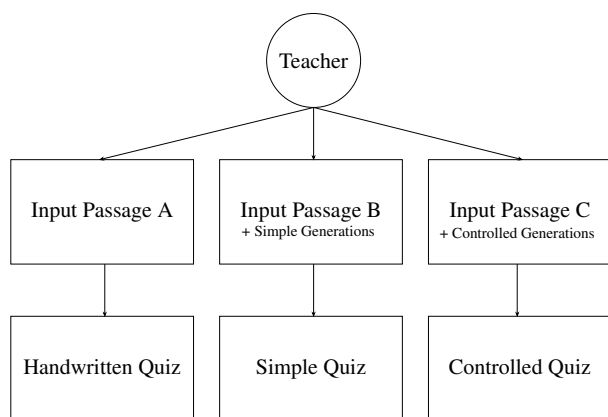


Figure 3: Quiz writing experiment diagram depicting the three quiz writing settings each teacher completed.

The teachers were asked to record their screen during the quiz writing process. These videos were analyzed to assess teachers’ experiences writing each kind of quiz. The *time* taken to write the quiz was measured, including how long the reading of the passage and additional candidates took. The *length* of the final quiz was measured, with a minimum of five and a maximum of 10 questions. And finally, the *source* of the questions was noted. In other words, it was recorded

if the question was directly copied from the GPT-3.5 generations, if it was copied and altered by the teacher, or if it was written completely from scratch.

Upon completion of the experiment, the teachers were asked to complete a brief post-quiz to better understand their feedback on the three quiz writing tasks. There was a section for free-form comments about the experiment. Then, the teachers were asked to pick which of the three quiz types was their preferred type and provide their reasoning.

Finally, eight other annotators (four per domain) analyzed the quality of the resulting quizzes using the metrics introduced in Section . They were recruited in the same way as outlined above, and have similar teaching experience and English proficiency. Six quizzes were seen by all four of the annotators in each domain in order to measure inter-annotator agreement. The rest of the quizzes were seen by two annotators. Having more than one annotator evaluate each quiz enables a more robust measurement of the quizzes’ quality. The annotators were not informed which quizzes included automatically generated questions.

Results

This section covers the key results for the evaluations of quiz quality and the quiz writing process. First, Section will compare and contrast the measured quality of the three different quiz types. Then, Section will discuss the teacher’s experiences in the different quiz writing settings.

Quiz Quality

Overall, the results from the quiz quality evaluations show that there is no notable loss of quality between the *handwritten* and the two generation quiz types. There is even an argument that an increase in quiz quality can be seen in some of the following results.

The three question-level metrics demonstrate a low count of *irrelevant*, *disfluent*, and *unanswerable* questions across all cohorts and quiz types. When comparing across quiz types, these results are relatively consistent. This implies that the use of generated candidates does not significantly increase or decrease the quality of quiz questions along these three aspects. The annotator agreement on these metrics is ‘fair’, with an average pairwise Cohen’s κ value of 0.3 for the BIO cohort and 0.6 for the ML cohort (Landis and Koch 1977).⁴

The *coverage* results show significant improvement when teachers utilize generated questions during their quiz writing process.⁵ Figure 4 depicts the increasing ratio of input text that is covered between *handwritten* and *simple*, as well as *simple* and *controlled* quiz types. In the ML cohort, a significant difference is observed between the *coverage* ratios

⁴Due to low percentages of *irrelevant*, *disfluent*, and *unanswerable* questions the agreement values are calculated on a unbalanced dataset and may not be an accurate representation of the real agreement on problem cases.

⁵N.b., a single *coverage* value for each quiz was annotated by the first author of this paper, so no measurement of agreement is available for this metric. We leave this for future work.

Metric	Handwritten		Simple		Controlled	
	BIO	ML	BIO	ML	BIO	ML
Irrelevant	0.03	0.00	0.05	0.05	0.04	0.00
Disfluent	0.14	0.11	0.09	0.08	0.14	0.11
Unanswerable	0.08	0.10	0.10	0.08	0.12	0.10

Table 1: Question-level quality metrics by cohort and quiz type. The values are the proportion of questions that are *irrelevant*, *disfluent*, or *unanswerable*.

in the *handwritten* and *controlled* quizzes. In the BIO cohort, we observe the same significant difference, as well as the difference between the *coverage* ratios in the *simple* and *controlled* quizzes. While it is not the only important aspect of a reading quiz, the increased *coverage* provided by automatically generated questions can benefit teachers in their writing process.




Figure 4: Context coverage by cohort and quiz type. * represents a significant difference at the $\alpha = 0.05$ level, and *** represents a significant difference at the $\alpha = 0.001$ level. The error bars represent 95% confidence intervals.

The three ordinal quiz-level metrics similarly have positive results across the board, as seen in the mean column of Table 2. The annotator agreement values are quite high. These metric’s agreement is measured with Kendall’s τ as it is suitable for ordinal scales (Schaeffer and Levitt 1956). All average pairwise values are above 0.5, and most above 0.81. Figure 5 shows the difference in quiz-level metrics between

Metric	Mean		Kendall’s τ	
	BIO	ML	BIO	ML
Structure [1,3]	2.35	2.43	0.81	0.83
Redundancy [1,3]	2.55	2.80	0.59	0.80
Usefulness [1,4]	3.12	2.97	0.95	0.52

Table 2: Quiz-level quality metrics and annotator agreement. The *structure* and *redundancy* metrics are on ordinal scales from 1 to 3, and the *usefulness* metric is on an ordinal scale from 1 to 4.

the three quiz types. Notably, for all three metrics and in both cohorts, one of the generation type quizzes is the highest

rated. In the ML cohort there is even a significant difference in the *usefulness* ratings between the *handwritten* and *simple* quizzes. This points to the fact that the generations can help improve the quality of quizzes. Further optimization of the generation process, perhaps with input from the teachers directly on what types of questions will be useful in a given setting, could show even stronger quality improvements.

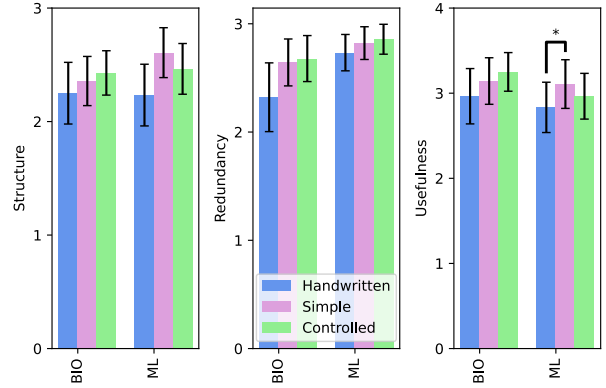


Figure 5: Quiz-level quality metrics. The error bars represent 95% confidence intervals. * represents a significant difference at the $\alpha = 0.05$ level.

Quiz Writing Experience

The quiz writing results demonstrate no consistent loss of efficiency when teachers use generated questions, at least with respect to the *time* taken. Figure 6 depicts the mean *time* to write each type of quiz for the two cohorts. The mean values are all relatively close to one another, and the 95% confidence intervals demonstrate that there appears to be no distinct difference in the quiz types with respect to *time*. Further analysis of these results demonstrate that the *time* metric is more dependent on the teacher and the particular passage than the generated candidate questions.




Figure 6: Time taken to create a quiz by cohort and quiz type. The error bars represent 95% confidence intervals.

The average *length* of the quizzes was 8.14 in the BIO cohort and 7.11 in the ML cohort. The mean number of questions in each quiz type was within one question to the mean across all types, demonstrating that the use of generations did not alter this variable. A more interesting comparison can be seen in Figure 7, where the *simple* and *controlled* quiz

types' question *sources* are compared. In both cohorts, the teachers hand write fewer questions when they have the *controlled* generations at their disposal. In fact, in the *ML* cohort, this difference is statistically significant. Again in both cohorts, teachers directly copy more of the *controlled* generations into their quizzes than the *simple* generations; and correspondingly they copy and edit fewer questions. This finding demonstrates that teachers choose to use the questions generated with Bloom's taxonomy more than other generations, providing motivation to continue QG research with pedagogical goals in mind.




Figure 7: Teachers use of generated questions by cohort. The error bars represent 95% confidence intervals. * represents a significant difference at the $\alpha = 0.05$ level.

The post-quiz results demonstrate a strong preference from the teachers in both cohorts towards the use of *controlled* generations, as depicted in Figure 8. The following comments from the teachers provide additional evidence for the fact that they find the *controlled* generations most useful for their quiz writing. Teachers stated that “*The generated questions for each specific type were incredibly useful.*” and “*I particularly liked the ‘creating’ questions, as I wouldn’t have come up with most of these myself.*” Another teacher commented that they “*tried to incorporate different command terms and level questions to prepare a quiz that tests the depth of a student’s understanding*”, indicating the intended usage of the questions in different levels of Bloom’s taxonomy. This finding corroborates this paper’s guiding intuition that QG methods should be designed with pedagogical goals in mind.




Figure 8: Teacher quiz type preference by cohort.

Limitations

It is important to note some limitations of this work. Firstly, we acknowledge that the quiz writing setting is to some extent contrived. In reality, teachers’ quiz writing experiences will be subjective: they might use additional resources or existing knowledge, write a draft and return to it later to edit, not constrain the number of questions, have different learning goals in mind, and more. The quiz writing setting was necessary to control certain variables for comparison within this work, but it creates the possibility that the results do not reflect the reality of quiz writing for teachers. Future work should aim to remove these constraints to better assess how teachers might realistically use generated questions. Secondly, this work only considers one LLM, two domains, the English-language setting, and a limited number of teachers. Although out of scope for this work, future work should aim to expand the variety of these aspects to assess how generated questions might benefit additional educational settings. Thirdly, a missing consideration in this work is the other half of the educational setting: students. Future work should include student goals, opinions, and performance in order to more comprehensively understand the implications of using automatically generated questions in the classroom. Despite these limitations, this work ultimately takes a step in the right direction for the future of QG research within realistic educational use cases.

Conclusion

This paper aims to show that LLMs are capable of generating different types of questions from a given context that teachers find useful to create a quiz that is of comparable quality to a handwritten version. To do this, quiz writing experiments were conducted comparing three types of quizzes: *handwritten*, *simple*, and *controlled* quizzes. The *controlled* quizzes utilized questions generated to correspond with the levels of Bloom’s taxonomy. The results demonstrate that teachers strongly prefer writing quizzes with the help of *controlled* generations. They also directly copy more of the *controlled* generations than the *simple* generations, indicating that these questions are of higher quality or better suited to a teacher’s goals. This confirms our hypothesis that teachers find automatically generated pedagogical questions useful for quiz writing. Additionally, an evaluation of quiz quality showed that the quizzes with both *controlled* and *simple* generations are of comparable quality. Some metrics even point towards their superior quality, when compared to *handwritten* quizzes. We hope that these findings will help to direct the future of educational QG research towards practical applications that meet the goals and needs of teachers and students.

Acknowledgements

The authors would like to thank Mitacs for their grant for this project, and CIFAR for their continued support. Additional thanks are extended to the teachers and annotators for their time and effort and the anonymous reviewers for their valuable feedback.

References

- Elkins, S.; Kochmar, E.; Serban, I.; and Cheung, J. C. 2023. How Useful Are Educational Questions Generated by Large Language Models? In *International Conference on Artificial Intelligence in Education*, 536–542. Springer.
- Kasneci, E.; Seßler, K.; Küchemann, S.; Bannert, M.; Dementieva, D.; Fischer, F.; Gasser, U.; Groh, G.; Günnemann, S.; Hüllermeier, E.; et al. 2023. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103: 102274.
- Krathwohl, D. R. 2002. A revision of Bloom’s taxonomy: An overview. *Theory into practice*, 41(4): 212–218.
- Kurdi, G.; Leo, J.; Parsia, B.; Sattler, U.; and Al-Emari, S. 2020. A systematic review of automatic question generation for educational purposes. *International Journal of Artificial Intelligence in Education*, 30: 121–204.
- Laban, P.; Wu, C.-S.; Murakhovs’ ka, L.; Liu, W.; and Xiong, C. 2022. Quiz design task: Helping teachers create quizzes with automated question generation. *arXiv preprint arXiv:2205.01730*.
- Landis, J. R.; and Koch, G. G. 1977. An application of hierarchical kappa-type statistics in the assessment of majority agreement among multiple observers. *Biometrics*, 363–374.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35.
- Mazidi, K.; and Nielsen, R. D. 2014. Pedagogical evaluation of automatically generated questions. In *Intelligent Tutoring Systems: 12th International Conference, ITS 2014, Honolulu, HI, USA, June 5-9, 2014. Proceedings 12*, 294–299. Springer.
- Mulla, N.; and Gharpure, P. 2023. Automatic question generation: a review of methodologies, datasets, evaluation metrics, and applications. *Progress in Artificial Intelligence*, 12(1): 1–32.
- Nenkova, A.; and Passonneau, R. J. 2004. Evaluating content selection in summarization: The pyramid method. In *Proceedings of the human language technology conference of the north american chapter of the association for computational linguistics: Hlt-naacl 2004*, 145–152.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback.
- Schaeffer, M. S.; and Levitt, E. E. 1956. Concerning Kendall’s tau, a nonparametric correlation coefficient. *Psychological Bulletin*, 53(4): 338.
- Steuer, T.; Bongard, L.; Uhlig, J.; and Zimmer, G. 2021. On the linguistic and pedagogical quality of automatic question generation via neural machine translation. In *Technology-Enhanced Learning for a Free, Safe, and Sustainable World: 16th European Conference on Technology Enhanced Learning, EC-TEL 2021, Bolzano, Italy, September 20-24, 2021, Proceedings 16*, 289–294. Springer.
- Van Campenhout, R.; Hubertz, M.; and Johnson, B. G. 2022. Evaluating AI-Generated Questions: A Mixed-Methods Analysis Using Question Data and Student Perceptions. In *Artificial Intelligence in Education: 23rd International Conference, AIED 2022, Durham, UK, July 27–31, 2022, Proceedings, Part I*, 344–353. Springer.
- Wang, X.; Fan, S.; Houghton, J.; and Wang, L. 2022a. Towards process-oriented, modular, and versatile question generation that meets educational needs. *arXiv preprint arXiv:2205.00355*.
- Wang, Z.; Valdez, J.; Basu Mallick, D.; and Baraniuk, R. G. 2022b. Towards human-like educational question generation with large language models. In *International conference on artificial intelligence in education*, 153–166. Springer.
- Zhang, H.; Song, H.; Li, S.; Zhou, M.; and Song, D. 2022. A survey of controllable text generation using transformer-based pre-trained language models. *arXiv preprint arXiv:2201.05337*.