

# FairytaleQA Translated: Enabling Educational Question and Answer Generation in Less-Resourced Languages<sup>\*</sup>

Bernardo Leite<sup>[0000-0002-9054-9501]</sup>, Tomás Freitas Osório<sup>[0009-0001-2036-3197]</sup>,  
and Henrique Lopes Cardoso<sup>[0000-0003-1252-7515]</sup>

Artificial Intelligence and Computer Science Laboratory (LIACC)  
Faculty of Engineering of the University of Porto (FEUP)  
Rua Dr. Roberto Frias, 4200-465 Porto, Portugal  
`bernardo.leite@fe.up.pt`  
`tomas.s.osorio@gmail.com`  
`hlc@fe.up.pt`

**Abstract.** Question Answering (QA) datasets are crucial in assessing reading comprehension skills for both machines and humans. While numerous datasets have been developed in English for this purpose, a noticeable void exists in less-resourced languages. To alleviate this gap, our paper introduces machine-translated versions of FairytaleQA, a renowned QA dataset designed to assess and enhance narrative comprehension skills in young children. By employing fine-tuned, modest-scale models, we establish benchmarks for both Question Generation (QG) and QA tasks within the translated datasets. In addition, we present a case study proposing a model for generating question-answer pairs, with an evaluation incorporating quality metrics such as question well-formedness, answerability, relevance, and children suitability. Our evaluation prioritizes quantifying and describing error cases, along with providing directions for future work. This paper contributes to the advancement of QA and QG research in less-resourced languages, promoting accessibility and inclusivity in the development of these models for reading comprehension. The code and data is publicly available at [github.com/bernardoleite/fairytaleqa-translated](https://github.com/bernardoleite/fairytaleqa-translated).

**Keywords:** Question Answering · Question Generation · Machine Translation.

## 1 Introduction

Question Answering (QA)<sup>1</sup> datasets, typically composed of text passages and corresponding question-answer pairs, have been used in advancing Question Generation (QG) and QA tasks research. Building these datasets often requires extensive effort from domain experts who are well-versed in the specific domain

<sup>\*</sup> Preprint - Accepted for publication at ECTEL 2024.

<sup>1</sup> Henceforth, we use “QA dataset” to denote datasets comprising questions and answers, and “QA task” to refer to the computational task of answering a question.

being targeted. A notable example is the FairytaleQA [22], a high-quality QA dataset carefully designed by education experts who applied evidence-based narrative comprehension frameworks [16,1] to annotate each question. While FairytaleQA excels in its quality and design, its content is exclusively tailored to a single language — English. This is also the case in other QA datasets, which presents a challenge for researchers aiming to extend their research to other languages. Notably, this challenge applies to languages that, despite not being low-resourced, lack QA datasets of this nature (referred to here as *less*-resourced languages). One viable solution involves relying on machine translation to generate reliable datasets, namely those that support the evaluation of neural models in downstream tasks. While machine translation is still not without issues, the resulting datasets can serve the purpose of fostering research in less-resourced languages.

Motivated by this observation, the first goal of this work is to provide machine-translated versions of FairytaleQA, in three<sup>2</sup> of the most widely spoken Romance languages: Spanish, Portuguese (European pt-PT and Brazilian pt-BR), and French. To estimate the quality of the translated versions, we conduct an error analysis of one of the translated datasets. As our second goal, we establish baseline benchmarks for both QA and QG tasks, achieved through fine-tuning modest-scale models. These models offer the advantage of being usable on more affordable hardware, in contrast to large language models that might incur costs for using expensive hardware or external APIs. Figure 1 shows an example where the original text and a QA pair are translated into Portuguese and applied in both QA and QG tasks.

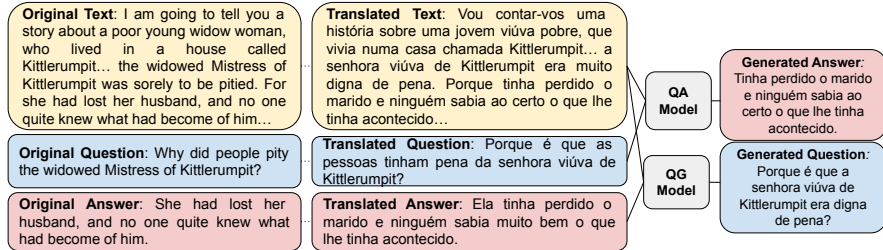


Fig. 1: QA and QG tasks using a translated text and corresponding QA pair.

The third goal entails our case study, which focuses on presenting and evaluating a model designed for Question-Answer Pair Generation (QAPG), specifically targeting the Portuguese language (pt-PT). Our motivation for generating both questions and answers is driven by the need for automated assessment of student answers in real-world educational scenarios. We aim particularly to understand

<sup>2</sup> We also have included translated datasets for Italian and Romanian in our repository, although they were not studied in this research.

the faithfulness of QA pairs generated by a modest-scale QAPG, trained on translated data, and compare them to those found in actual exams. This forms the central research question of our case study:

RQ: *Can a modest-scale QAPG model, trained on translated data, generate QA pairs that are qualitatively similar to those used in real exams in a less-resourced language?*

We conduct human evaluation based on criteria such as well-formedness, answerability, relevance, and suitability for children. While our case study focuses on a single language, we expect it can serve as (1) a guideline for researchers seeking to explore QA/QG in other translated datasets, and (2) a representative illustration of expected errors during QA pair generation, along with directions for future work.

Our analysis found that the QAPG model trained on translated data can generate well-formed questions. However, detection of semantic ambiguities was noted. Additionally, our observations revealed that the alignment of the generated QA pairs is not always consistent, indicating the need for further refinement. In summary, our contributions are:

- We provide machine-translated versions<sup>3</sup> of the FairytaleQA dataset in three widely spoken Romance languages.
- We introduce fine-tuned models<sup>4</sup> along with benchmarks for both QA and QG tasks within the translated datasets.
- We present a case study where a QAPG model is proposed and evaluated, focusing on quantifying and describing expected errors in QA pair generation.

## 2 Background and Related Work

### 2.1 Question & Answer Generation Corpora

Over the years, numerous QA and QG datasets<sup>5</sup> have been proposed. A reference example is the SQuAD dataset [18], comprising crowd-sourced QA pairs extracted from Wikipedia articles. In the case of SQuAD, to assess its viability in other languages, some machine-translated versions have been made publicly available, including Spanish [5] and Slovak [19]. However, a limitation of these datasets is that they are built upon open-domain and general-purpose resources like Wikipedia and news articles. As a result, models trained on this type of data (or translations) may generate suitable questions/answers for a generic domain but are likely to fall short in serving specific educational purposes.

Focusing on corpora explicitly created for an educational purpose, RACE [10], CLOTH [21], and FairytaleQA [22] are prime examples of this accomplishment. The questions within RACE and CLOTH were derived from English exams

<sup>3</sup> Made available at [github.com/bernardoleite/fairytaleqa-translated](https://github.com/bernardoleite/fairytaleqa-translated).

<sup>4</sup> Made available at [github.com/bernardoleite/fairytaleqa-translated](https://github.com/bernardoleite/fairytaleqa-translated).

<sup>5</sup> QA datasets are typically also used for QG.

aimed at middle to high school students. Notably, CLOTH labels each question with its corresponding reasoning level, providing a valuable advantage for educational applications.

**Why FairytaleQA?** We selected the FairytaleQA dataset due to its texts and corresponding QA pairs aligning with the goal of *supporting narrative comprehension*. As highlighted by Xu *et al.* [22], narrative comprehension is a high-level skill that strongly correlates with reading success [15]. Furthermore, narrative stories possess a well-defined structure comprising distinct elements and their relationships. We also selected this dataset due to its structured focus on specific dimensions of reading comprehension skills. Specifically, education experts have annotated each question, following evidence-based narrative comprehension frameworks [16,1], and addressing two key attributes: *narrative elements* and *explicitness* (explained in Section 3.1). To the best of our knowledge, this is the first work involving translations of FairytaleQA.

## 2.2 Question & Answer Generation Tasks

While the QA task revolves around finding or generating an answer for a certain question, QG involves the automatic generation of well-structured questions. Both tasks can leverage diverse data sources, including free text or knowledge bases. QG holds practical significance in education. It can assist teachers in creating assessments or stimulate students’ self-learning [9].

From a methodological standpoint, prior work on QA and QG has primarily focused on fine-tuning small to medium-sized pre-trained language models [13,23,8] for generating questions and answers. With the emergence of large language models [3], recent research leveraged the prompting paradigm, where a textual query specifies the desired generation task given a set of prompt examples. Although this technique has demonstrated compelling performance [7], it can be reliant on hardware or external API access with associated costs, posing a limitation. In this study, we concentrate on employing modest-scale fine-tuned models. The term “modest-scale” is used because we specifically employ T5-based models [17] with parameters in the magnitude of millions, making it suitable for more affordable hardware. This is in contrast to large language models, which typically have parameters in the billions<sup>6</sup>. Nevertheless, we compare the performance of the case study’s model (T5) against a large language model, i.e., GPT-4 (Section 5).

## 3 Translating the FairytaleQA Dataset

### 3.1 About FairytaleQA

FairytaleQA includes 10,580 QA pairs, manually created by educational experts based on 278 stories. On average, each story comprises 15 sections, and each

<sup>6</sup> It is worth noting recent efforts to make these models more accessible, for instance through quantization techniques.

section (paragraph), made of multiple sentences, contains an average of 3 questions. Each QA pair is annotated with corresponding narrative elements and explicitness labels. Narrative elements are described as follows:

- **Character:** These require test takers to identify or describe the characteristics of story characters.
- **Setting:** These inquire about the place or time where story events occur and typically start with “Where” or “When”.
- **Action:** These focus on characters’ behaviors or seek information about their actions.
- **Feeling:** These explore characters’ emotional status or reactions, often framed as “How did/does/do...feel”.
- **Causal relationship:** These examine the cause-and-effect relationships between two events, often starting with “Why” or “What made/makes”.
- **Outcome resolution:** These ask for the events that result from prior actions in the story, typically phrased as “What happened/happens/has happened...after...”.
- **Prediction:** These request predictions about the unknown outcome of a particular event based on existing information in the text.

Question explicitness is defined as follows:

- **Explicit:** The answers are directly present in the text and can be located within specific passages.
- **Implicit:** Answers cannot be directly pinpointed in the text, requiring the ability to summarize and make inferences based on implicit information.

### 3.2 Machine Translation and Sample Evaluation

Each dataset instance contains a text section (paragraph) and a QA pair. Consequently, to translate each instance effectively, we needed to address these three components. We opted to translate all three components together, separated by a line break. This approach was chosen based on empirical evidence from preliminary test translations, indicating that translating questions and answers with improved contextualization of the section text led to greater coherence between these components. For machine translation, we have used DeepL<sup>7</sup>, a commercial tool known for its translation quality compared to its competitors. We translated the original English version into the three most spoken Romance languages: Spanish, Portuguese (European pt-PT and Brazilian pt-BR), and French.

To assess the translation quality of the data, we conducted a manual error analysis on 10 translated texts and corresponding 150 QA pairs (we selected the sections with the highest number of QA pairs). Our analysis focused on the translated version in European Portuguese (pt-PT). As all translated versions are Romance languages, any identified error types are expected to be potentially applicable to these versions, serving as a cautionary note for researchers

<sup>7</sup> <https://www.deepl.com/translator>

considering their utilization. The quantification of issues (by text, question, and answer) is presented in Table 1 and elaborated upon below.

Table 1: Incidence of issues resulting from machine translation in a dataset sample.

Issue	Nr. Texts	Nr. Questions	Nr. Answers
Translating Names	7/10	22/150	10/150
Change of Gender	0/10	1/150	2/150
Lost in Translation	0/10	1/150	0/150
Outdated Spelling Agreement	1/10	1/150	0/150

**Translating Names:** A common issue arises from the non-deterministic translation of proper or common names. Take, for instance, the example “*Tom* had a little more trouble with him.” which translates to “O *Tomás* teve um pouco mais de problemas com ele.” This undesired translation of the character’s name can have implications, especially when the name is translated in the question/answer but not in the text, introducing inconsistency between components. A solution to this issue could involve employing an alignment approach using Named Entity Recognition models.

**Change of Gender:** Some cases have been detected where the translation implies a gender change. For instance, consider the QA pair: “How will the mother feel when the horses keep getting cut into two? Angry.” which translates to: “Como se sentirá a mãe quando os cavalos continuarem a ser cortados em dois? Zangado.” In the translation, the answer *Angry* is rendered as *Zangado* which is masculine, instead of *Zangada*, which is the appropriate feminine form referring to *mãe* (*mother*). While finding a straightforward solution to this issue is not apparent, one option could be adjusting gender terms using set linguistic rules.

**Lost in Translation:** We detected one case where a question’s meaning changed during translation. This occurred when translating “What was the chariot drawn by?” which, in back-translation, resulted in “What chariot was pulling?”. We attribute this to the translator’s performance, and at the moment, a solution is not apparent for us.

**Outdated Spelling Agreement:** We detected one word that did not comply with the latest spelling agreement. This occurrence may vary depending on the timeliness of the data on which the translator has been created. The solution might be using regular expressions, for example, to replace the most recent terms using a local dictionary. Additionally, spell-checking could serve as a viable option.

## 4 Baseline Benchmarks: Question & Answer Generation

In this section, we present benchmarks for both Question Answering (QA) and Question Generation (QG) tasks within FairytaleQA.

### 4.1 Implementation Details

Leveraging pre-trained T5 [17] encoder-decoder models, known for their remarkable performance in QA/QG tasks [20], we employ the *t5-base*<sup>8</sup> version pre-trained on each language’s data (monolingual). Figure 2 shows the setup for the QA and QG tasks. For QA, the encoder concatenates the question and text, and the decoder generates the answer. In QG, the encoder concatenates the answer and text, and the decoder generates the question. We use special labels to differentiate the components. Our maximum token input is set to 512, while the maximum token output is set to 128. During training, the models undergo a maximum of 20 epochs and incorporate early stopping with a patience of 2. A batch size of 16 is employed. During inference, we utilize beam search with a beam width of 5. We use the original train/val/test splits that contain 8,548/1,025/1,007 QA pairs.

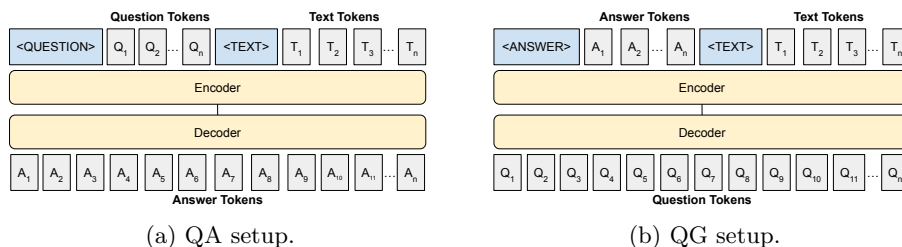


Fig. 2: Setups for Question Answering (QA) and Question Generation (QG).

### 4.2 Results

Table 2 displays the results obtained on the test set, reported using the  $\text{ROUGE}_L$ -F1 metric employed by the FairytaleQA authors. It computes the  $n$ -gram similarity (0-1) between the ground-truth questions/answers and the generated questions/answers. Our overall results for QG in the original English language slightly exceed those reported by the authors of FairytaleQA (0.527 vs. 0.530), as do for overall QA (0.536 vs. 0.551). This indicates that our experimental setup aligns with the expected outcomes. It should be noted that the values obtained

<sup>8</sup> <https://huggingface.co/t5-base>

<sup>9</sup> This model and consequent results are those reported by the authors of FairytaleQA.

Table 2: QA and QG benchmarks according to the language of the translated dataset. Values are obtained through ROUGE<sub>L</sub>-F1 metric (0-1).

Language	Model	QG	QA		
		Overall	Overall	Explicit	Implicit
English	BART <sup>9</sup> [14]	0.527	0.536	0.620	0.286
English	T5 [17]	0.530	0.551	0.679	0.167
Spanish	T5S [2]	0.445	0.382	0.474	0.111
Portuguese (pt-PT)	PTT5 [4]	0.496	0.436	0.520	0.185
Portuguese (pt-BR)	PTT5 [4]	0.470	0.448	0.534	0.192
French	T5-fr [6]	0.404	0.431	0.518	0.172

in QA (Implicit) consistently appear lower than those for QA (Explicit). This discrepancy is consistent with the observations made by FairytaleQA’s authors, as questions requiring implicit answers necessitate inference or summarization from the source text.

For QG, all the translated versions produced lower values than the English baseline, with statistically significant differences (t-test with  $p < .05$ ). There is also statistical significance when comparing only among the translated versions, although the effect size is not large (Cohen’s  $d < 0.44$ ). For the QA task, we observe similar trends. However, the results of QA reveal promising evidence when delving into the explicitness aspect. As seen in the original English version, there is a substantial difference between explicit and implicit QA pairs, indicating that the translated QA pairs maintain this consistency in explicitness.

To assess the impact of translation on model performance, we decided to back-translate the Spanish dataset, which exhibited the poorest QA performance, into English. The newly translated version in English achieved the following ROUGE<sub>L</sub>-F1 results: 0.497 (QG) and 0.478 (QA, overall). Notably, the QA result surpassed the performance of previous translated versions and approximated the level of the original English version<sup>10</sup>. Hence, we refrain from attributing the lower results solely to the quality of the translated versions. It is plausible that the monolingual models themselves may exert an influence.

## 5 Case Study: Presenting and Evaluating a QAPG Model

In an educational context, for effective assessment, AI systems must provide not only questions but also their corresponding answer pairs. Therefore, we introduce a Question-Answer Pair Generation (QAPG) model. In this case study, our goal is to assess the extent to which QA pairs, generated using a modest-scale QAPG model and trained on translated data, can approximate real exam QA pairs and those generated using a large language model (i.e., GPT-4). The case study focuses on European Portuguese (pt-PT). While not empirically demonstrated,

<sup>10</sup> While the difference is statistically significant, the effect size is small (Cohen’s  $d < 0.19$ ).

we anticipate that the insights gained from the model evaluation will likely have broader applicability to other languages.

### 5.1 Introducing QAPG

We frame QAPG as a controllable QA pair generation task, where narrative labels (recall Section 3.1) are used as control attributes. Figure 3 illustrates the QAPG model setup for generating a QA pair targeting the *action* narrative element. The control attribute  $\langle \text{NAR} \rangle$  is added at the start of the input, preceding the section text. The decoder outputs labels that differentiate the  $\langle \text{QUESTION} \rangle$  and  $\langle \text{ANSWER} \rangle$  components of the output. This is based on and supported by recent studies [8,12,11], where the idea is to guide the model to generate a QA pair of the intended type. Indeed, the controllable process is reported to hold significant importance in the educational field [9], as it enhances the creation of customized questions tailored to a student’s specific needs. In our case study, we explore using different control attributes to generate diverse and multiple QA pairs from the same text. Implementation details are similar to those explained in Section 4.1.

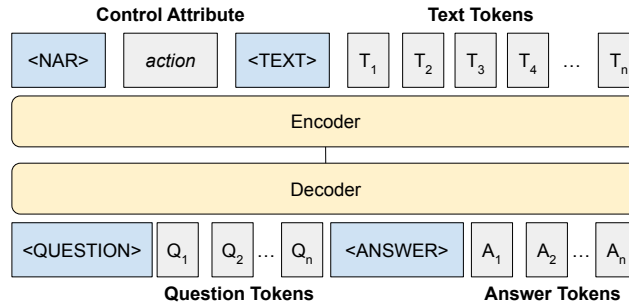


Fig. 3: Setup for Question-Answer Pair Generation (QAPG), based on [11].

### 5.2 Evaluation Protocol

To evaluate the quality of the generated QA pairs, we engaged a group of 15 volunteers who are native speakers of Portuguese with a higher education degree. Each individual was instructed to read a real-exam narrative text and rate a set of 15 mixed QA pairs about that text. The set included 5 QA pairs directly extracted from a real exam, 5 generated by the QAPG model<sup>11</sup>, and 5 generated by the GPT-4<sup>12</sup> model. We conducted 4 sets of inquiries, totaling 60 QA

<sup>11</sup> The 5 questions were randomly selected from a pool of 7, corresponding to the 7 narrative elements that can be controlled.

<sup>12</sup> We prompt GPT-4 Turbo to generate QA pairs from the narrative text in the target language, tailored for children aged 7 to 10, aligned with the FairytaleQA’s audience.

pairs evaluated. The texts were sourced from real exams provided by the Portuguese Ministry of Education<sup>13</sup> for elementary school students, aligning with FairytaleQA’s target audience. For each question, participants were instructed to assess it based on the following metrics:

- **Well-formedness** (question only): Is this question well-formed? [Yes: well-formed with no errors; No: ill-formed with orthographic, grammatical or semantic errors].
- **Relevance with Text** (question only): Is the question grounded in the text? [Yes; No].
- **Children Suitability** (question only): Do you consider the question suitable for a child between 7 and 10 years old? [5-point Likert scale].
- **Answerability** (question only): Is there an answer to the question in the text? [Yes; No].
- **Answer Alignment** (QA pair): Examine this potential answer: GENERATED ANSWER (*the evaluators have seen the generated answer*). Is this answer valid for the question? [Yes; No].
- **Observations** (optional): Please justify/comment on your rating.

We ensured that at least 3 participants evaluated each QA pair.

### 5.3 Results

Table 3 displays the results, with each cell represented as XX/YY, where XX is the count of QA pairs with a majority positive voting result, and YY with a majority negative voting result<sup>14</sup>. An exception is the suitability metric, which is an overall average. The subsequent analysis focuses on identifying questions where, for certain metrics, the resulting vote was negative. We also emphasize cases where a positive vote result did not exhibit unanimity (100% agreement).

Table 3: Human evaluation results, determined by majority vote, except for the suitability metric, which represents an overall average.

Metric	QA Pairs Provenance		
	Real-Exam	GPT-4	QAPG
Well-formedness	20/0	19/1	20/0
Relevance with Text	20/0	20/0	19/1
Answerability	20/0	20/0	14/6
Answer Alignment	18/2	20/0	8/12
Children Suitability	4,77	4,83	4,68

For **well-formedness**, the questions generated by QAPG were all rated positively, showcasing the model’s ability to produce questions without orthographic

<sup>13</sup> <https://iave.pt/provas-e-exames/provas-e-exames/provas-de-afizacao-eb/>

<sup>14</sup> In 4% of the cases, votes were tied, leading us to seek an additional volunteer to break the tie.

or grammar errors. Two positive cases did not exhibit unanimity. Evaluators who disagreed offered valuable feedback, suggesting changes in the main verb tense for enhanced clarity. For example<sup>15</sup>, “What *will* the boy ask...?” can be improved to “What *did* the boy ask...?”. For GPT-4, one question was considered ill-formed, and another lacked unanimity. Both were identified as written in a different variant: pt-BR (instead of pt-PT). This limitation in GPT-4 motivates further exploration, as our prompt specifically included the instruction to generate QA pairs in the intended language variant.

For **relevance with text**, despite the overall positive score, 1 question generated by QAPG received a negative vote. This question inquired about the feelings mistakenly attributed to a wrong character in the story. There were 3 questions lacking unanimity. In the first two cases, evaluators pointed out semantic inconsistencies in the context of the story, as seen in “Who paid for the hard-boiled egg?”, which cannot be resolved based on the story. Additionally, one question was raised regarding its lack of coherence with the chronological order of events described in the text.

For **answerability**, 6 questions generated by QAPG were voted negatively, indicating that, according to the evaluators, the generated questions have no answer in the text. Upon analysis, we identified that 1 question was flagged since the required answer pertains to an event that is unclear in the story. Additionally, 2 questions were noted to demand answers that are not explicitly provided in the text. Finally, 3 questions were reported to be fully unanswerable in the text, such as, “Who was the bear?”, a character whose description is missing in the text.

For **answer alignment**, which assesses whether the generated answer is valid for the question, QAPG shows more pronounced negative results. The first 6 negative cases are the ones already spotted for lack of answerability. In the remaining 6 cases, the generated answer was reported as inaccurate (2 cases), incomplete (1 case), or wrong/nonsensical (3 cases). Surprisingly, 2 real-exam answers were evaluated as incorrect. Upon analysis, we detected that these answers concern a temporal location, i.e., “The basket appeared a few days later.” and a spatial location, “The bear lives in the middle of nowhere.” which are not very specific but not necessarily wrong.

Overall, the evaluators considered the questions to be **suitable** ( $\geq 4.6$ ) for a child aged between 7 and 10.

## 6 Discussion and Future Work

We revisit our research question: *Can a modest-scale QAPG model, trained on translated data, generate QA pairs that are qualitatively similar to those used in real exams in a less-resourced language?*

We found that the QAPG model, trained on translated data, is capable of generating well-formulated QA pairs. Although some suggestions for improvement were noted, no orthographic or grammatical errors were spotted. This is

<sup>15</sup> The following examples have been translated from Portuguese to English.

encouraging, particularly given that the dataset on which the model was trained was obtained through machine translation. Consequently, the QAPG’s ability to generate well-formed questions is on par with those from real exams and GPT-4. Then, we found that the QAPG model can generate questions that are relevant to the text. However, 1 question was reported as irrelevant, and 3 questions showed a high potential for issues related to coherence and semantics. Consequently, QAPG’s ability to generate relevant questions is estimated to be slightly lower than that of questions from real exams and GPT-4. We also discovered that 14/20 questions generated by the QAPG model were reported as answerable. Consequently, the QAPG’s ability to generate answerable questions is estimated to be also slightly lower than that of questions from real exams and GPT-4. Ultimately, we identified the primary challenge of QAPG to be the generation of aligned QA pairs. Only 8/20 of the generated answers can be considered valid to the corresponding questions. This marks a significant difference when compared to the alignment observed in real exams and GPT-4. Consequently, the QAPG’s ability to generate aligned QA pairs is estimated to be significantly lower than that of questions from real exams and GPT-4. Here, we summarize the main issues observed when applying our QAPG model:

- Improper verb tense in the generated question;
- Incoherent or semantically inconsistent generated question within the context of the story;
- Generated question lacks an answer within the text;
- Generated question is not aligned with the generated answer.

Overall, the results represent a preliminary yet encouraging demonstration of the potential of QA/QG systems trained on translated educational data for less-resourced languages. Despite the progress, numerous challenges persist. Considering the above points, we identify the following areas for future work as promising:

- **Exploring alternative modest-scale models:** As we observed in Section 4, the performance of the monolingual QA/QG model likely influences the results. Considering recent efforts to enhance accessibility to large language models through techniques like quantization, we consider exploring alternative recent models to assess their impact on results.
- **Double-checking answer existence:** After generating the QA pairs, some questions were identified as having no answer in the text. One approach to confirming answer existence is by applying an additional and robust question-answering model. Comparing the answer from this model with the initially generated answer in the QA pair may provide a strong indication of answer presence.
- **Employing QA pair alignment verification models:** The primary challenge identified in our results lies in generating QA pairs that exhibit alignment. We suggest exploring classification models capable of classifying QA pairs as aligned or not. This approach can facilitate the filtering of unwanted QA pairs, a procedure that we find lacking in the existing QA/QG literature. Double-checking answer existence can also be useful here.

- **Comparing between synthetic and translated data:** GPT-4 exhibited notable superiority over the modest-scale model trained on translated data. To explore potential data-related problems, we consider conducting an experiment: employ a large language model to generate QA pairs (via zero or few-shot prompting) in the target language and then train modest-scale models on these QA pairs. This allows for a comparison between models trained on synthetic data and those trained on translated QA pairs.

## 7 Conclusion

In this work, we have advanced the accessibility of machine-translated versions of the FairytaleQA dataset, which is useful for assessing reading comprehension skills. Additionally, we developed and analyzed baseline QA/QG models. In our case study, focusing on a model trained to generate QA pairs from translated data, we successfully demonstrated the generation of well-formulated questions. However, the primary challenge lies in ensuring effective alignment within QA pairs. Although our work is limited to analyses in a single language, we hope our insights will prove valuable for researchers aiming to replicate and enhance our methods in alternative languages. For immediate future work, we intend to develop a QA pair alignment verification model to filter unwanted QA pairs.

## Acknowledgments

The authors thank the 15 participants who voluntarily participated in the evaluation process. We would like to specifically acknowledge those who agreed to share their names: Ângela Esteves, Bruno Miguel Pinto, David Reis, Diana Pinto, Maria de Abreu, Mariana Coelho, Pedro Costa, Rui Leixo and Vítor Magalhães. This work was financially supported by Base Funding - UIDB/00027/2020 and Programatic Funding - UIDP/00027/2020 of the Artificial Intelligence and Computer Science Laboratory – LIACC - funded by national funds through the FCT/MCTES (PIDDAC). Bernardo Leite is supported by a PhD studentship (with reference 2021.05432.BD), funded by Fundação para a Ciência e a Tecnologia (FCT).

## References

1. Alonzo, J., Basaraba, D., Tindal, G., Carriveau, R.S.: They read, but how well do they understand? an empirical look at the nuances of measuring reading comprehension. *Assessment for Effective Intervention* **35**(1), 34–44 (2009)
2. Araujo, V., Trusca, M.M., Tufiño, R., Moens, M.F.: Sequence-to-sequence spanish pre-trained language models. arXiv preprint arXiv:2309.11259 (2023), model URL: [huggingface.co/vgaraujov/t5-base-spanish](https://huggingface.co/vgaraujov/t5-base-spanish)
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C.,

- Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 1877–1901. Curran Associates, Inc. (2020), [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf)
4. Carmo, D., Piau, M., Campiotti, I., Nogueira, R., Lotufo, R.: Ptt5: Pre-training and validating the t5 model on brazilian portuguese data. *arXiv preprint arXiv:2008.09144* (2020), model URL: [huggingface.co/unicamp-dl/ptt5-base-portuguese-vocab](https://huggingface.co/unicamp-dl/ptt5-base-portuguese-vocab)
  5. Carrino, C.P., Costa-jussà, M.R., Fonollosa, J.A.R.: Automatic Spanish translation of SQuAD dataset for multi-lingual question answering. In: *Proceedings of the 12th Language Resources and Evaluation Conference*. pp. 5515–5523. European Language Resources Association, Marseille, France (May 2020), <https://aclanthology.org/2020.lrec-1.677>
  6. Dublineau, J.: T5 question generation and question answering. <https://huggingface.co/JDBN/t5-base-fr-qg-fquad> (2020)
  7. Elkins, S., Kochmar, E., Serban, I., Cheung, J.C.K.: How useful are educational questions generated by large language models? In: Wang, N., Rebolledo-Mendez, G., Dimitrova, V., Matsuda, N., Santos, O.C. (eds.) *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky*. pp. 536–542. Springer Nature Switzerland, Cham (2023)
  8. Ghanem, B., Lutz Coleman, L., Rivard Dexter, J., von der Ohe, S., Fyshe, A.: Question generation for reading comprehension assessment by modeling how and what to ask. In: *Findings of the Association for Computational Linguistics: ACL 2022*. pp. 2131–2146. Association for Computational Linguistics, Dublin, Ireland (May 2022). <https://doi.org/10.18653/v1/2022.findings-acl.168>, <https://aclanthology.org/2022.findings-acl.168>
  9. Kurdi, G., Leo, J., Parsia, B., Sattler, U., Al-Emari, S.: A systematic review of automatic question generation for educational purposes. *International Journal of Artificial Intelligence in Education* **30**(1), 121–204 (2020)
  10. Lai, G., Xie, Q., Liu, H., Yang, Y., Hovy, E.: RACE: Large-scale ReAding comprehension dataset from examinations. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pp. 785–794. Association for Computational Linguistics, Copenhagen, Denmark (Sep 2017). <https://doi.org/10.18653/v1/D17-1082>, <https://aclanthology.org/D17-1082>
  11. Leite, B., Cardoso, H.: On few-shot prompting for controllable question-answer generation in narrative comprehension. In: *Proceedings of the 16th International Conference on Computer Supported Education - Volume 2: CSEDU*. pp. 63–74. INSTICC, SciTePress (2024). <https://doi.org/10.5220/0012623800003693>
  12. Leite, B., Cardoso, H.L.: Towards enriched controllability for educational question generation. In: Wang, N., Rebolledo-Mendez, G., Matsuda, N., Santos, O.C., Dimitrova, V. (eds.) *Artificial Intelligence in Education*. pp. 786–791. Springer Nature Switzerland, Cham (2023)
  13. Leite, B., Lopes Cardoso, H.: Neural question generation for the portuguese language: A preliminary study. In: Marreiros, G., Martins, B., Paiva, A., Ribeiro, B., Sardinha, A. (eds.) *Progress in Artificial Intelligence*. pp. 780–793. Springer International Publishing, Cham (2022)

14. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: Jurafsky, D., Chai, J., Schluter, N., Tetreault, J. (eds.) *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. pp. 7871–7880. Association for Computational Linguistics, Online (Jul 2020). <https://doi.org/10.18653/v1/2020.acl-main.703>, <https://aclanthology.org/2020.acl-main.703>
15. Lynch, J.S., Van Den Broek, P., Kremer, K.E., Kendeou, P., White, M.J., Lorch, E.P.: The development of narrative comprehension and its relation to other early reading skills. *Reading Psychology* **29**(4), 327–365 (2008)
16. Paris, A.H., Paris, S.G.: Assessing narrative comprehension in young children. *Reading Research Quarterly* **38**(1), 36–76 (2003)
17. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* **21**(140), 1–67 (2020), <http://jmlr.org/papers/v21/20-074.html>, model URL: [huggingface.co/google-t5/t5-base](https://huggingface.co/google-t5/t5-base)
18. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: SQuAD: 100,000+ questions for machine comprehension of text. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. pp. 2383–2392. ACL, Austin, Texas (Nov 2016). <https://doi.org/10.18653/v1/D16-1264>, <https://aclanthology.org/D16-1264>
19. Staš, J., Hládek, D., Kocúr, T.: Slovak question answering dataset based on the machine translation of the squad v2. 0. *Journal of Linguistics/Jazykovedný casopis* **74**(1), 381–390 (2023)
20. Ushio, A., Alva-Manchego, F., Camacho-Collados, J.: Generative language models for paragraph-level question generation. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. pp. 670–688. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Dec 2022), <https://aclanthology.org/2022.emnlp-main.42>
21. Xie, Q., Lai, G., Dai, Z., Hovy, E.: Large-scale cloze test dataset created by teachers. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. pp. 2344–2356. Association for Computational Linguistics, Brussels, Belgium (Oct-Nov 2018). <https://doi.org/10.18653/v1/D18-1257>, <https://aclanthology.org/D18-1257>
22. Xu, Y., Wang, D., Yu, M., Ritchie, D., Yao, B., Wu, T., Zhang, Z., Li, T., Bradford, N., Sun, B., Hoang, T., Sang, Y., Hou, Y., Ma, X., Yang, D., Peng, N., Yu, Z., Warschauer, M.: Fantastic questions and where to find them: FairytaleQA – an authentic dataset for narrative comprehension. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 447–460. Association for Computational Linguistics, Dublin, Ireland (May 2022). <https://doi.org/10.18653/v1/2022.acl-long.34>
23. Zhao, Z., Hou, Y., Wang, D., Yu, M., Liu, C., Ma, X.: Educational question generation of children storybooks via question type distribution learning and event-centric summarization. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. pp. 5073–5085. Association for Computational Linguistics, Dublin, Ireland (May 2022). <https://doi.org/10.18653/v1/2022.acl-long.348>, <https://aclanthology.org/2022.acl-long.348>