

Generative Language Models for Paragraph-Level Question Generation

Asahi Ushio and Fernando Alva-Manchego and Jose Camacho-Collados
Cardiff NLP, School of Computer Science and Informatics, Cardiff University, UK
{UshioA,AlvaManchegoF,CamachoColladosJ}@cardiff.ac.uk

Abstract

Powerful generative models have led to recent progress in *question generation* (QG). However, it is difficult to measure advances in QG research since there are no standardized resources that allow a uniform comparison among approaches. In this paper, we introduce QG-Bench, a multilingual and multidomain benchmark for QG that unifies existing question answering datasets by converting them to a standard QG setting. It includes general-purpose datasets such as SQuAD (Rajpurkar et al., 2016) for English, datasets from ten domains and two styles, as well as datasets in eight different languages. Using QG-Bench as a reference, we perform an extensive analysis of the capabilities of language models for the task. First, we propose robust QG baselines based on fine-tuning generative language models. Then, we complement automatic evaluation based on standard metrics with an extensive manual evaluation, which in turn sheds light on the difficulty of evaluating QG models. Finally, we analyse both the domain adaptability of these models as well as the effectiveness of multilingual models in languages other than English. QG-Bench is released along with the fine-tuned models presented in the paper,¹ which are also available as a demo.²

1 Introduction

Question generation (QG, Mitkov and Ha, 2003) is the task of generating a question given an input context consisting of a document, a paragraph or a sentence, and an answer where the question is anchored (see Figure 1). QG has been widely studied in natural language processing communities (Du et al., 2017; Zhou et al., 2017; Du and Cardie, 2018), and it has recently been exploited to train question answering (QA) models without human supervision (Lewis et al., 2019; Zhang and

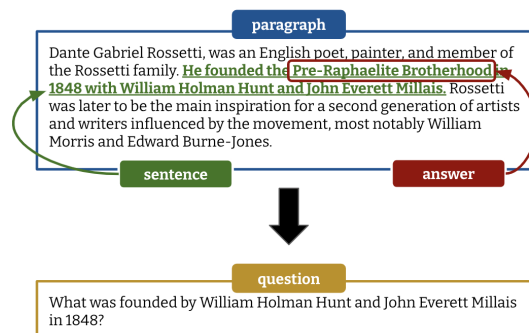


Figure 1: Overview of paragraph-level QG.

Bansal, 2019; Puri et al., 2020), or as a means of data augmentation (Shakeri et al., 2020; Bartolo et al., 2021). It has also been applied to develop educational systems (Heilman and Smith, 2010; Lindberg et al., 2013), information retrieval models (Pyatkin et al., 2021; Lewis et al., 2021), and for model interpretation (Perez et al., 2020; Lee et al., 2020).

Despite its success in downstream applications, the development of neural QG models has received less attention. For example, the choice of the base pre-trained model is arbitrary (without proper justification in most cases) as it is not straightforward to compare different models. As a consequence, while ERNIE-GEN (Xiao et al., 2021) and UniLMv2 (Bao et al., 2020) are current SotA in the SQuAD QG benchmark (Du et al., 2017), T5 (Raffel et al., 2020) and BART (Lewis et al., 2020a) are used in many applications in practice (Paranjape et al., 2021; Bartolo et al., 2021; Lewis et al., 2021; Pyatkin et al., 2021).

A possible reason is inconsistent evaluation and comparison of QG models, due to the lack of appropriate evaluation protocols and benchmarks. For instance, evaluation of QG models relies on BLEU4 (Papineni et al., 2002), METEOR (Denkowski and Lavie, 2014), and ROUGE_L (Lin, 2004), with human-made questions as references. However,

¹<https://github.com/asahi417/lm-question-generation>

²<https://autoqg.net/>

some of these metrics may have low correlation with human judgements, especially when it comes to *answerability*, since they tend not to take the associated answer into account (Nema and Khapra, 2018). Moreover, QG applications can use different contexts as input, such as sentence-level (Pyatkin et al., 2021; Lewis et al., 2019) vs paragraph-level (Zhang and Bansal, 2019; Puri et al., 2020), or answer-aware (Shakeri et al., 2020; Bartolo et al., 2021) vs answer-free (Lopez et al., 2020). These are generally used interchangeably in the literature.

To investigate how to tackle the issues previously raised, we introduce QG-Bench, a collection of standard QA datasets unified into a single benchmark, including domain-specific datasets and for eight different languages (§ 3). We then use QG-Bench to fine-tune various generative language models (LMs) by formulating paragraph-level QG as a sequence-to-sequence generation task (§ 4), and measure their performance on in-domain and language-specific data (§ 5). Finally, we present a multi-faceted analysis of our QG models by varying their input context size (§ 6.1), conducting a manual evaluation (§ 6.2), and studying their abilities for domain adaptation (§ 6.3).

2 Related Work

Early work on QG was based on human-engineered templates (Mitkov and Ha, 2003; Rus et al., 2010) and well-designed pipelines (Heilman and Smith, 2010; Labutov et al., 2015), but soon neural approaches took over by generating a question from a text in an end-to-end manner (Du et al., 2017; Zhou et al., 2017; Du and Cardie, 2018). The quality of QG models was later improved by masked LM pre-training (Devlin et al., 2019; Liu et al., 2019) where the encoder of the QG model is fine-tuned from pre-trained LMs (Chan and Fan, 2019; Zhang and Bansal, 2019). Recently, sequence-to-sequence LM pre-training has allowed to fully fine-tune QG models (both encoder and decoder), achieving SotA performance (Dong et al., 2019; Qi et al., 2020; Bao et al., 2020; Xiao et al., 2021). Following the latest research in the literature, we focus on sequence-to-sequence LM-based QG models.

QG can be applied to domain adaptation (Shakeri et al., 2020), knowledge-enhanced LM pre-training (Jia et al., 2021), adversarial/counterfactual data augmentation (Bartolo et al., 2021; Paranjape et al., 2021), and nearest neighbour QA systems (Lewis et al., 2021). Applications

of QG go beyond QA, including semantic role labeling (Pyatkin et al., 2021), visual QA (Krishna et al., 2019), multi-hop question decomposition (Perez et al., 2020), and question rewriting (Lee et al., 2020). Moreover, QG can be applied to unsupervised QA, which consists of training a QA model without any supervision and relying on a QG model to generate questions (Lewis et al., 2019). Puri et al. (2020) showed that with a carefully-designed QG model, we can generate high-quality QA datasets on which a QA model can even outperform their supervised counterparts. This inspired Zhang and Bansal (2019) to propose QA-based evaluation, which connects the quality of a QG model to the accuracy of a QA model trained on the synthetic data generated by the QG model.

While QG models can be applied to this variety of tasks, the comparison across tasks is not always straightforward. For this reason, and given the relevance of QG in current research, in this paper we propose an intrinsic QG benchmark in which we can evaluate different aspects of a QG model in a simple manner, including, but not only, analysis of input types, domain adaptability and multilinguality. The most similar work to ours is the MTG benchmark (Chen et al., 2021), which contains multilingual test sets for four NLG tasks. While QG is part of this benchmark, there are a few major differences from our proposed QG-Bench: (i) we provide training/validation/test sets to allow model training in each language in addition to the evaluation; (ii) MTG’s test set consists of parallel sentences across languages by a translation from English, while we leverage monolingual datasets; (iii) we include eight languages, while MTG has five; and (iv) QG-Bench includes datasets from different domains and styles.

3 QG-Bench: A Unified Question Generation Benchmark

In this section, we describe our process to construct QG-Bench, including data collection and unification (§ 3.1), and its statistics (§ 3.2).

3.1 Data Collection and Unification

We unified a collection of datasets, designed to be used for QG model training and evaluation. All datasets are in the same format, where each entry contains four features: *paragraph*, *sentence*, *question*, and *answer*. As described in Figure 1, we assume *question* as the output of a QG system,

which is conditioned by an *answer* and it is always a sub-string of a *sentence* from a *paragraph*. We leverage existing QA datasets by compiling them into this unified QG format. All datasets included in QG-Bench are described below.

SQuAD (English). We first consider SQuAD v1.1 (Rajpurkar et al., 2016), an extractive QA dataset based on Wikipedia which has been used in QG commonly since (Du et al., 2017; Zhou et al., 2017). As the original test set of SQuAD is not released, we use the same data split as in (Du et al., 2017).

Domain-specific Datasets (English). To assess models’ domain adaptivity, we consider two domain-specific QA datasets: SQuADShifts (Miller et al., 2020) and SubjQA (Bjerva et al., 2020). SQuADShifts contains questions in the same style of SQuAD but from four additional domains (Amazon/Wikipedia/News/Reddit), while SubjQA consists, unlike SQuAD, of subjective questions/answers in general (e.g. *how is the hero? - the hero was wonderful*) across six domains. As the original SQuADShifts consists of test set only, we created a new training/validation/test split, in which half of the dataset remains in the test set, while the remaining half is split for validation and training by a 1:2 ratio.

Datasets in Languages other than English. To investigate multilinguality in QG, we compile the following seven SQuAD-style QA datasets: JAQuAD (So et al., 2022) (Japanese), GerQuAD (Möller et al., 2021) (German), SberQuAd (Efimov et al., 2020) (Russian), KorQuAD (Lim et al., 2019) (Korean), FQuAD (d’Hoffschmidt et al., 2020) (French), Spanish SQuAD (Casimiro Pio et al., 2019) (Spanish), and Italian SQuAD (Croce et al., 2018) (Italian). Since they do not release test sets, we sampled a subset from the training sets as the test set following Du et al. (2017). The test sets contain the same number of questions as its validation set, and the new training/test splits have no overlap in terms of the paragraphs.

Other Datasets not Included in QG-Bench. In theory, any extractive QA dataset could be part of our benchmark. However, we decided not to include datasets such as BioASQ (Tsatsaronis et al., 2015) and NewsQA (Trischler et al., 2017) because they have very long input texts, representing another category that needs extra mechanisms to handle long sequences (Izcard and Grave, 2020a,b), which is out of the scope of this paper. In addition, one could leverage multilingual QA benchmarks

	Data size (train/valid/test)	Average character length (para./sent./ques./ans.)
SQuAD	75,722 / 10,570 / 11,877	757 / 179 / 59 / 20
SubjQA		
- <i>Book</i>	637 / 92 / 191	1,514 / 146 / 28 / 83
- <i>Elec.</i>	697 / 99 / 238	1,282 / 129 / 26 / 66
- <i>Grocery</i>	687 / 101 / 379	896 / 107 / 25 / 49
- <i>Movie</i>	724 / 101 / 154	1,746 / 146 / 27 / 72
- <i>Rest.</i>	823 / 129 / 136	1,006 / 104 / 26 / 51
- <i>Trip</i>	875 / 143 / 397	1,002 / 108 / 27 / 51
SQuADShifts		
- <i>Amazon</i>	3,295 / 1,648 / 4,942	773 / 111 / 43 / 18
- <i>Wiki</i>	2,646 / 1,323 / 3,969	773 / 184 / 58 / 26
- <i>News</i>	3,355 / 1,678 / 5,032	781 / 169 / 51 / 20
- <i>Reddit</i>	3,268 / 1,634 / 4,901	774 / 116 / 45 / 19
Multilingual QG		
- <i>Ja</i>	27,809 / 3,939 / 3,939	424 / 72 / 32 / 6
- <i>Es</i>	77,025 / 10,570 / 10,570	781 / 122 / 64 / 21
- <i>De</i>	9,314 / 2,204 / 2,204	1,577 / 165 / 59 / 66
- <i>Ru</i>	40,291 / 5,036 / 5,036	754 / 174 / 64 / 26
- <i>Ko</i>	54,556 / 5,766 / 5,766	521 / 81 / 34 / 6
- <i>It</i>	46,550 / 7,609 / 7,609	807 / 124 / 66 / 16
- <i>Fr</i>	17,543 / 3,188 / 3,188	797 / 160 / 57 / 23

Table 1: Statistics of all datasets integrating into our question generation benchmark after unification.

(Clark et al., 2020; Artetxe et al., 2020; Lewis et al., 2020b) to obtain multilingual QG datasets, but XQuAD (Artetxe et al., 2020) and MLQA (Lewis et al., 2020b) do not contain training sets, and TydiQA (Clark et al., 2020) contains a very small training set. Instead, we focused on monolingual QA datasets in each language.

3.2 Data Statistics

Table 1 summarizes statistics of each QG dataset after unification. It can be observed that SubjQA and SQuADShifts have ten to a hundred times less training data than SQuAD. Also, SubjQA’s answers are twice longer than SQuAD’s answers, which can be explained by how they differ in the way questions are formed (i.e., SubjQA being more subjective in nature). Likewise, except for Spanish, the datasets for languages other than English contain less training data than the original SQuAD, with the number varying depending on the language.

4 LMs for Question Generation

In this section, we formalize the QG task from a language modelling perspective (§ 4.1), including details on the fine-tuning process (§ 4.2) and the setup for our experiments with QG-Bench (§ 4.3).

4.1 Task Formulation

Given an input text x , the goal of QG is to generate a natural question \hat{q} related to the information in

the input. The task is formulated as a conditional sequence generation, and the model is optimized to maximize the conditional log-likelihood $P(q|x)$ as in Equation 1.

$$\hat{q} = \arg \max_q P(q|x) \quad (1)$$

In practice, the log-likelihood is factorized into word or subword level predictions, similar to other sequence-to-sequence learning settings (Sutskever et al., 2014).

4.2 Language Model Fine-tuning

Fine-tuning sequence-to-sequence LMs on QG can be done in the same way as for Machine Translation or Summarization, where models are trained to predict the output tokens given the input tokens (Dong et al., 2019; Qi et al., 2020; Bao et al., 2020; Xiao et al., 2021). We follow Chan and Fan (2019) by introducing a highlight token `<hl>` to take into account an answer a within a context c as below:

$$x = [c_1, \dots, \text{<hl>}, a_1, \dots, a_{|a|}, \text{<hl>}, \dots, c_{|c|}]$$

Instead of a paragraph, we can similarly use a sentence to highlight an answer (sentence-level QG) or highlight a sentence instead of an answer (answer-free QG). We investigate these model variations in our analysis (§ 6.1), but assume the answer highlighted paragraph as the default input.

Note that it is possible to train other types of LMs on QG, but masked LMs were not designed for natural language generation and require a specific decoding technique (Chan and Fan, 2019). Also, recurrent LMs have poor ability for conditional generation on the answer due to its unidirectional architecture (Lopez et al., 2020). Since they are not as suited for QG as the sequence-to-sequence models, they are out of the scope of this paper.

4.3 Experimental Setup

Comparison Models. As sequence-to-sequence LMs, we use T5 (Raffel et al., 2020) and BART (Lewis et al., 2020a) for the English datasets and mT5 (Xue et al., 2021) and mBART (Liu et al., 2020) for the multilingual experiments. Model weights are taken from HuggingFace (Wolf et al., 2020).³ Previous research reported improvements on QG with more recent LMs (Qi et al., 2020;

³We use t5-small, t5-base, t5-large, facebook/bart-base, facebook/bart-large, and google/mt5-small.

Xiao et al., 2021; Bao et al., 2020). We tried to replicate these previous works in QG-Bench, but after multiple attempts using their provided code and contacting the authors, this was not possible. Nonetheless, both T5 and BART are widely used in practice and, as we will show, they can still provide strong results with an appropriate configuration.

Parameter Optimization. We performed an extensive exploration to find the best combination of hyper-parameters to fine-tune LMs on QG, which consists of a two-phase search. First, we fine-tune a model on every possible configuration from the search space for 2 epochs. The top-5 models in terms of BLEU4 (Papineni et al., 2002) on the validation set are selected to continue fine-tuning until their performance plateaus.⁴ Finally, the model that achieves the highest BLEU4 on the validation set is employed as the final model. We used BLEU4 as an objective metric in our parameter optimization since it is light to compute, and following previous work (Du and Cardie, 2018; Dong et al., 2019; Xiao et al., 2021). However, as we will see in our experiments, future work could also explore the usage of alternative metrics for validation. The search space contains 24 configurations, which are made up of learning rates from [0.0001, 0.00005, 0.00001], label smoothing from [0.0, 0.15], and batch size from [64, 128, 256, 512].⁵ Our experiments show that this simple parameter optimization strategy significantly improves all models’ performances by robustly finding the best configuration for each one.⁶

We ran the parameter optimization on a machine equipped with two Nvidia Quadro RTX 8000. Taking SQuAD as a reference, training and evaluation took around three weeks for T5_{LARGE}, one week for T5_{BASE} and mT5_{SMALL}, three days for T5_{SMALL}, one week for BART_{LARGE}, and four days for BART_{SMALL}.

5 Automatic Evaluation

In this section, we report the main results in QG-Bench (§ 3), using the methodology described in § 4.

⁴This two-stage process is introduced due to computation limitations, and we might see further improvements (even if small) if a full validation search is performed.

⁵Other parameters are fixed: random seed is 1, beam size is 4, input token length is 512, and output token length is 34 for fine-tuning and 64 for evaluation.

⁶See Appendix for the actual parameters found by the optimization procedure as well as more training details.

Model	Param	B4	R-L	MTR	BS	MS
NQG (Du et al.)	30M	12.28	39.75	16.62	-	-
UniLM (Dong et al.)	340M	22.78	51.57	25.49	-	-
UniLMv2 (Bao et al.)	110M	24.70	52.13	26.33	-	-
ProphetNet (Qi et al.)	340M	23.91	52.26	26.60	-	-
ERNIE-G (Xiao et al.)	340M	25.40	52.84	26.92	-	-
BART _{BASE}	140M	24.68	52.66	26.05	90.87	64.47
BART _{LARGE}	400M	26.17	53.85	27.07	91.00	64.99
T5 _{SMALL}	60M	24.40	51.43	25.84	90.45	63.89
T5 _{BASE}	220M	26.13	53.33	26.97	90.84	64.74
T5 _{LARGE}	770M	27.21	54.13	27.70	91.00	65.29

Table 2: QG model fine-tuning results on the test set of SQuAD where the best result in each metric is in bold face. The results in the top row group are existing SotA models taken from their original papers, while the bottom row contains our models.

5.1 Evaluation Metrics

To evaluate QG models, BLEU4 (B4, Papineni et al., 2002), METEOR (MTR, Denkowski and Lavie, 2014), and ROUGE_L (R-L, Lin, 2004) are commonly used to compare the generated outputs against reference questions at sentence level. We also compute BERTScore (BS, Zhang et al., 2019) and MoverScore (MS, Zhao et al., 2019). Both leverage BERT-like models on their computation, achieving higher correlations with human judgments than other traditional metrics in various NLG tasks (Zhang et al., 2019; Zhao et al., 2019). To the best of our knowledge, they have not been applied in QG evaluation before, regardless of their success in NLG. We use the default configuration for both metrics, which make use of RoBERTa_{LARGE} (Liu et al., 2019) for BERTScore and DistilBERT_{BASE} (Sanh et al., 2019) for MoverScore.

5.2 Results

SQuAD. Table 2 shows our results on the SQuAD test set along with other reported results from the literature. T5_{LARGE} provides the best results overall according to all automatic metrics. Even parameter-efficient models such as T5_{BASE} outperform ERNIE-GEN (Xiao et al., 2021), and T5_{SMALL} performs competitively with UniLMv2 (Bao et al., 2020) with nearly half the parameters. UniLMv2, in particular, was proposed as a highly-effective model in spite of its light weight. According to these results, T5_{SMALL} is also competitive on the QG task while being significantly lighter than other models. While T5 attains the best overall results, BART also proves competitive. In fact, BART_{BASE} is slightly better than T5_{BASE}

	Model	B4	R-L	MTR	BS	MS
English	mT5 _{SMALL}	21.65	48.95	23.83	90.01	62.75
	mT5 _{BASE}	23.03	50.67	25.18	90.23	63.60
	mBART	23.03	50.58	25.10	90.36	63.63
Russian	mT5 _{SMALL}	16.31	31.39	26.39	84.27	62.49
	mT5 _{BASE}	17.63	33.02	28.48	85.82	64.56
	mBART	18.80	34.18	29.30	87.18	65.88
Japanese	mT5 _{SMALL}	30.49	50.88	29.03	80.87	58.67
	mT5 _{BASE}	32.54	52.67	30.58	81.77	59.68
	mBART	32.16	52.95	29.97	82.26	59.88
Italian	mT5 _{SMALL}	7.37	21.93	17.57	80.80	56.79
	mT5 _{BASE}	7.70	22.51	18.00	81.16	57.11
	mBART	7.13	21.69	17.97	80.63	56.84
Korean	mT5 _{SMALL}	10.57	25.64	27.52	82.89	82.49
	mT5 _{BASE}	12.18	28.57	29.62	84.52	83.36
	mBART	10.92	27.76	30.23	83.89	82.95
Spanish	mT5 _{SMALL}	9.61	24.62	22.71	84.07	59.06
	mT5 _{BASE}	10.15	25.45	23.43	84.47	59.62
	mBART	9.18	24.26	22.95	83.58	58.91
German	mT5 _{SMALL}	0.43	10.08	11.47	79.90	54.64
	mT5 _{BASE}	0.87	11.10	13.65	80.39	55.73
	mBART	0.75	11.19	13.71	80.77	55.88
French	mT5 _{SMALL}	8.55	28.56	17.51	80.71	56.50
	mT5 _{BASE}	6.14	25.88	15.55	77.81	54.58
	mBART	0.72	16.40	7.78	71.48	50.35

Table 3: QG model fine-tuning results on the test set of all language-specific QG-Bench datasets where the best result in each language is in bold face.

and BART_{LARGE} is equal to T5_{LARGE} according to BERTScore. In general, it is hard to reliably compare different model architectures for the QG tasks, as there are different possible confounding factors including the model size. To have a more complete picture on the final performance, we complement this initial automatic evaluation in SQuAD with an extensive manual evaluation in § 6.2.

Language-specific Datasets. Table 3 presents the results on each language-specific dataset in QG-Bench with mT5_{SMALL}, mT5_{BASE}, and mBART. As this work introduces the first ever comprehensive multilingual QG model training/evaluation, these results can be viewed as baselines for future work in multilingual QG. Compared to results in English SQuAD, scores in multilingual QG are mostly worse than the smallest English model (T5_{SMALL}), which showcases the difficulties of non-English QG. Some languages are notably underperforming, which can be partially explained by the size of the training set. As we see in § 3.2, some datasets such as German and French have a limited amount of training instances, resulting in underfitting models for those languages. In general, the low scores in non-English datasets can be attributed to the under-

Domain	Model	B4	R-L	MTR	BS	MS
Amazon	BART _{BASE}	9.92	27.94	22.78	92.77	63.25
	BART _{LARGE}	9.80	28.69	23.79	92.49	63.31
	T5 _{SMALL}	8.41	27.04	22.17	91.89	62.11
	T5 _{BASE}	9.80	28.94	23.85	92.43	63.27
	T5 _{LARGE}	10.42	29.51	24.39	92.65	63.71
Wiki	BART _{BASE}	11.50	29.00	26.60	93.12	65.86
	BART _{LARGE}	12.12	29.94	27.12	93.39	66.22
	T5 _{SMALL}	10.90	28.18	25.95	92.63	65.04
	T5 _{BASE}	11.67	29.49	27.04	93.07	65.94
	T5 _{LARGE}	12.04	30.10	27.67	93.13	66.31
News	BART _{BASE}	8.78	24.85	25.13	92.86	64.99
	BART _{LARGE}	8.74	25.28	25.08	93.04	65.02
	T5 _{SMALL}	7.71	23.43	23.70	92.20	63.71
	T5 _{BASE}	8.53	24.93	25.21	92.68	64.70
	T5 _{LARGE}	9.16	25.97	25.98	93.01	65.46
Reddit	BART _{BASE}	8.78	26.03	22.57	92.32	62.35
	BART _{LARGE}	9.31	27.31	23.75	92.50	62.64
	T5 _{SMALL}	7.60	24.90	21.90	91.70	61.39
	T5 _{BASE}	8.75	26.84	23.57	92.26	62.52
	T5 _{LARGE}	9.16	27.24	23.97	92.43	62.74
Book	BART _{BASE}	2.03	23.24	20.57	92.96	62.85
	BART _{LARGE}	0.00	23.71	20.6	92.84	62.45
	T5 _{SMALL}	0.00	19.77	18.52	92.40	61.46
	T5 _{BASE}	0.00	22.95	21.20	93.32	63.14
	T5 _{LARGE}	0.00	23.68	20.83	92.89	62.51
Elec.	BART _{BASE}	3.83	29.41	25.08	93.76	66.00
	BART _{LARGE}	5.18	28.87	25.17	93.51	65.68
	T5 _{SMALL}	0.00	29.65	26.95	94.18	68.29
	T5 _{BASE}	4.55	29.99	27.39	94.26	68.33
	T5 _{LARGE}	4.57	30.55	27.56	94.27	68.80
Grocery	BART _{BASE}	1.82	24.54	20.8	94.09	65.76
	BART _{LARGE}	1.93	24.28	20.42	94.1	65.79
	T5 _{SMALL}	0.00	22.17	23.31	93.24	65.64
	T5 _{BASE}	0.83	15.63	19.87	90.56	61.47
	T5 _{LARGE}	1.13	17.40	20.64	91.39	63.41
Movie	BART _{BASE}	3.89	25.43	20.55	93.61	62.91
	BART _{LARGE}	4.21	25.92	21.64	93.23	62.4
	T5 _{SMALL}	0.00	25.76	22.54	94.08	64.63
	T5 _{BASE}	2.65	26.33	23.11	94.13	64.91
	T5 _{LARGE}	0.00	25.06	21.70	93.64	63.88
Rest.	BART _{BASE}	3.43	24.26	21.35	93.23	62.67
	BART _{LARGE}	5.54	24.77	22.46	93.23	63.57
	T5 _{SMALL}	0.00	11.72	13.21	87.81	55.42
	T5 _{BASE}	0.00	11.96	14.75	88.48	56.19
	T5 _{LARGE}	4.19	24.94	21.99	93.22	63.25
Trip	BART _{BASE}	4.79	26.37	25.26	93.92	64.91
	BART _{LARGE}	5.66	26.5	24.32	93.85	64.02
	T5 _{SMALL}	2.49	23.91	25.56	93.75	66.57
	T5 _{BASE}	1.74	16.06	20.13	90.76	59.70
	T5 _{LARGE}	5.35	27.69	27.45	94.46	67.76

Table 4: QG model fine-tuning results on the test set of SQuADShifts and SubjQA where the best result in each metric is in bold face.

lying model, so scaling up the model could lead to better performances in future work.

Domain-specific Datasets. Table 4 shows the results from all domain-specific datasets included in QG-Bench: SQuADShifts and SubjQA. Since each domain contains a small training set, our main strategy to achieve domain-specific QG models is to

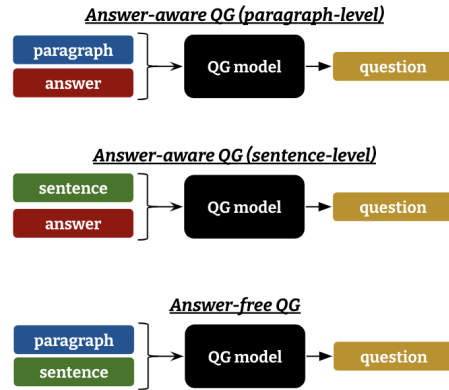


Figure 2: Input variations of QG models.

initialize their weights with a SQuAD fine-tuned model, and continue fine-tuning on the domain-specific training set (more details on different strategies in § 6.3). As expected, given the subjective nature of the dataset, results on SubjQA are generally low for most metrics, except for BERTScore whose score is even higher than in SQuAD in some cases. This implies that a model’s prediction may have less word-overlap against the true question, while its semantics is close to the true question to some extent.

6 Analysis

In this section, we complement the automatic evaluation with an extensive analysis on various relevant aspects of the question generation models.

6.1 Model Input

In our main experiments, the model input is the paragraph in which the answer is highlighted, as described in § 4.2. Here we explore variations of the QG model’s input type to understand the effect of different types of context. Concretely we consider two additional variants: *sentence-level* models that only take as input the sentence that contains the answer (instead of the whole paragraph); and *answer-free* models that highlight the sentence in the paragraph instead of the answer. Figure 2 provides a summary of the three different input types analysed.

In Table 5 we report automatic metrics from answer-free models and sentence-level QG models on SQuAD. In general, paragraph-based models, which use the most complete input, attain the best overall results. For example, answer-free T5_{LARGE} performs worse than paragraph-level T5_{SMALL} in all the metrics except METEOR, which indicates

Model	B4	R-L	MTR	BS	MS	
Answer-free	BART _{BASE}	21.97	49.70	23.72	90.38	63.07
	BART _{LARGE}	23.47	50.25	24.94	90.28	63.28
	T5 _{SMALL}	21.12	47.47	23.38	89.64	62.07
	T5 _{BASE}	22.86	49.51	24.52	90.03	62.99
	T5 _{LARGE}	24.27	51.30	25.67	90.41	63.97
Sent-level	BART _{BASE}	23.86	51.43	25.18	90.70	63.85
	BART _{LARGE}	23.86	51.43	25.18	90.70	63.85
	T5 _{SMALL}	23.23	50.18	24.80	90.36	63.18
	T5 _{BASE}	24.33	51.81	25.81	90.73	64.00
	T5 _{LARGE}	25.36	52.53	26.28	90.88	64.44
Para-level	BART _{BASE}	24.68	52.66	26.05	90.87	64.47
	BART _{LARGE}	26.17	53.85	27.07	91.00	64.99
	T5 _{SMALL}	24.40	51.43	25.84	90.45	63.89
	T5 _{BASE}	26.13	53.33	26.97	90.84	64.74
	T5 _{LARGE}	27.21	54.13	27.70	91.00	65.29

Table 5: QG model fine-tuning results on the test set of SQuAD for answer-free and sentence/paragraph-level QG models. The best overall result for each metric is in boldface.

the importance of the answer at question generation. Nonetheless, not having the answer as input provides competitive results, which may appear to be surprising given the incomplete input. When comparing sentence-level and paragraph-level, the difference is reduced, but paragraph-level models consistently outperform their sentence-level counterparts, even when smaller models are used. This implies that models actually utilize the global context provided by the full paragraph when it is available, rather than the more local information within the sentence only.

6.2 Manual Evaluation

Given the limitations of automatic metrics in text generation research (Reiter, 2018; Bhandari et al., 2020; Alva-Manchego et al., 2021), we also conducted a manual evaluation using Amazon Mechanical Turk, focusing on three criteria: *grammaticality* (i.e. grammatical correctness), *understandability* (i.e. whether the question is easy to be understood by readers) and *answerability* (i.e. whether the question can be answered by the given input answer).⁷ We randomly sampled 500 unique paragraphs from the SQuAD test set and selected a single answer in each paragraph. For each of the 500 paragraph-answer pairs, we generated questions from six QG models, and asked human anno-

⁷Understandability could correlate with grammaticality, but a question without any grammatical mistakes can have low understandability due to an over complex structure. Likewise, a question can be understandable even with a few grammatical mistakes. Annotation guidelines are included in the Appendix.

Model	Manual Metric			Automatic Metric				
	Ans.	Gra.	Und.	B4	R-L	MTR	BS	MS
NQG	1.21	2.35	2.63	3.33	14.30	33.53	88.27	58.25
BART _{LARGE}	2.70	2.89	2.93	16.15	29.93	51.35	90.95	65.44
T5 _{SMALL}	2.51	2.83	2.90	13.43	27.38	48.86	90.41	64.27
T5 _{LARGE}	2.80	2.93	2.95	17.56	30.42	52.00	90.94	66.09
- sent-level	2.47	2.91	2.95	14.88	27.49	48.97	90.76	64.53
- answer-free	2.46	2.91	2.95	13.62	26.82	47.37	90.20	64.00

Table 6: Manual evaluation results along with the automatic metrics. Each score is averaged within the 500 questions for the evaluation where the best result in each metric is in bold face.

tators to score them for the criteria with a 3-points scale. Each question was evaluated by five judges, thus collecting a total of 15,000 human judgments. As quality control, we asked workers to be native English speakers, and instructed them to do a qualification test first, and only those who passed the test worked on our annotation task. The given time of each assignment (with each assignment containing ten instances to annotate) was 30 minutes, and the reward of the annotation task was \$2 per assignment.⁸ We attach a screenshot of the annotation interface in the Appendix.

Comparison Models. For the manual evaluation, the target QG models include T5_{LARGE}, T5_{SMALL} and BART_{LARGE} based paragraph-level QG models; T5_{LARGE} sentence-level and answer-free QG models; and NQG (Du et al., 2017), which is based on an LSTM-architecture. NQG is included for completeness and to better analyse the effect of pre-trained LMs in general. T5_{LARGE} is our best model according to automatic metrics, so we compare it against different input types (answer-free or sentence-level), different sizes (T5_{SMALL}), and different model architectures (BART_{LARGE}).

Inter-annotator Agreement. Since there are five unique annotators per each generated question, we calculated Fleiss’s kappa to measure the inter-annotator-agreement. We obtained 0.30 and 0.36 for grammaticality and understandability respectively, resulting in fair-agreement (Landis and Koch, 1977). The kappa is 0.61 in answerability, which is a substantial-agreement.

Model-wise Evaluation. We report the results of our manual evaluation in Table 6, where each score is averaged over the 500 questions used in the study. Answerability is the most affected by model size/context and type/model architecture, compared to the other metrics, except for NQG, which is

⁸The full price of annotation exercise was about \$3,000.

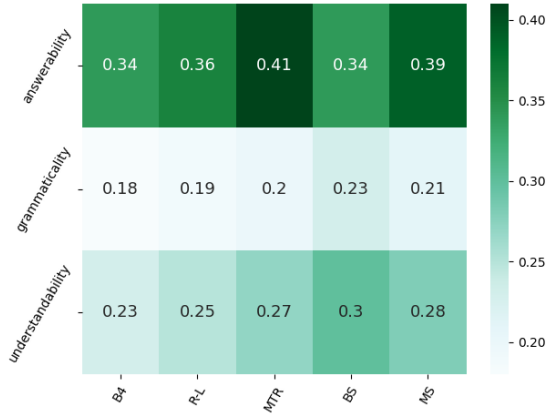


Figure 3: Spearman’s rank correlation over all the generated questions within the manual evaluation.

the only non-LM pre-training based approach. In fact, when we compare $T5_{LARGE}$ ’s paragraph-level against sentence-level, answerability decreases unlike the other two criteria, highlighting the importance of including all relevant context available so that the model can generate a suitable question. On the other hand, while answer-free models are worse than sentence-level models according to automatic metrics, the manual evaluation does not reflect a significant difference between them. In general, we can see how $T5_{LARGE}$, which is the best model overall according to the automatic metrics, is also the most robust model overall according to the manual evaluation, which reinforces the conclusions from the automatic evaluation.

Correlation Analysis. Leveraging the large dataset of collected human judgments, we investigate the correlation between human annotations and the automatic metrics considered in the automatic evaluation (§ 5.2). For this analysis, we included all the generated questions from all the models considered in the manual evaluation. This means 3,000 generated questions from six diverse models where each question receives five annotations. We took the average across all the five annotators for each generated question to compute the correlation. Figure 3 shows the Spearman’s rank correlation coefficient across the automatic metrics and the criteria collected through our manual evaluation.⁹ The p-values of all correlations are less than 0.05, so they are all statistically significant. To check the significance of the increase in the correlation across metrics, we ran a William test, showing that the

⁹See the full correlation analysis in Appendix § B.2.

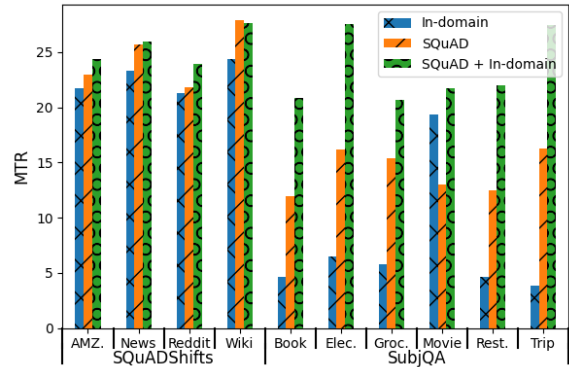


Figure 4: Comparison of METEOR (MTR) scores for $T5_{LARGE}$ across in-domain fine-tuning, zero-shot transfer of SQuAD fine-tuned model, and in-domain fine-tuning from SQuAD model.

increase is statistically significant in all cases.¹⁰

According to the correlation analysis, no metric achieved a high agreement with human judgements in all criteria. This means that we should not rely on a single metric to capture all quality aspects of a model’s output. We can conclude, however, that METEOR and MoverScore are well-aligned with human judgements on answerability, while BERTScore appears to be better suited for grammaticality and understandability. Most importantly, BLEU4 and ROUGE_L, which have been mostly used in the QG literature as default metrics, are not as reliable as the other metrics in any criteria.

6.3 Domain Adaptation

In our main experiments in the domain-specific datasets of QG-Bench (§ 5.2), models were initialized by the SQuAD fine-tuned model due to the limited training set in each domain. To further explore the domain adaptability of QG models, we compared three different setups: (1) fine-tuning in the in-domain training set without SQuAD initialization, (2) zero-shot transfer from the SQuAD fine-tuned model, and (3) fine-tuning with a prior SQuAD initialization. Figure 4 shows the results of $T5_{LARGE}$ (the best model in most of the domains in Table 4 and the manual evaluation) in each domain for those three settings. For this analysis, we focus on the METEOR metric, which attains the highest correlation with human judges in answerability.¹¹ We can confirm that the best setup is to initialize the model on SQuAD and then further fine-tune it

¹⁰Full results of the William test are in Appendix § B.3.

¹¹The full set of results for other metrics and models is available in the Appendix, with similar general trends.

on the domain-specific training sets. For SQuAD-Shifts, however, this improvement is less marked in general, suggesting that T5 can handle inputs from different domains to a certain extent. In contrast, the zero-shot setting with SQuAD fine-tuning in SubjQA achieves very poor results overall. This is to a certain extent expected since the questions in SubjQA are of very different styles.

Finally, while in this section we focused on the domain adaptability for English, in the Appendix we also show zero-shot cross-lingual transfer results, adapting English-training models to other languages. Similarly to previous work (Chen et al., 2021), the main conclusion is that there is still significant room for improvement for zero-shot cross-lingual transfer in QG.

7 Conclusion

In this paper we presented QG-Bench, a unified benchmark and evaluation for testing paragraph-level QG models. The benchmark is composed of the general-purpose SQuAD dataset, as well as domain-specific datasets of different styles for English. Moreover, it includes language-specific datasets for eight different languages. Using QG-Bench as a reference, we tested recent generative language models on the task, and evaluated them across a range of automatic metrics. To complement the automatic evaluation, we performed a comprehensive manual evaluation to better understand the performance of models and the role of automatic metrics (e.g., our study shows there are better metrics than the popular BLEU4 when it comes to QG). In general, our results show that LMs have come a long way for QG, being very competitive (e.g., T5 attains an overall manual score of, respectively, 2.80, 2.93 and 2.95 in answerability, grammaticality and understandability on SQuAD), but have room for improvement when dealing with different domains and styles, and especially on languages other than English.

As future work, we will continue to study QG evaluation metrics in-depth to better understand what aspects we are missing when we use specific automatic metrics, using our manual evaluation as a proxy. Moreover, the QG models analysed in this paper require an answer to be specified beforehand to generate the question. As a way to relax the constrain, we can train models for question and answer pair generation (QAG) by generating the answer together with the question given a context.

By generating both answers and questions together, new evaluation metrics would also be required to understand the validity and diversity of the answers selected, which we leave for future work.

Limitations

In this paper, we explored paragraph-level QG models, which limits their input up to around 500 tokens, and the same methodology cannot be easily applied to longer documents. In multilingual QG modeling, we considered datasets in seven different languages, but all of them are medium- to high-resource languages, so our experimental results cannot be generalized to a truly low-resource language setting. Finally, although the focus of our paper is mostly in SQuAD-style one hop extractive QA, QG is also studied in more complex scenarios such as multi-hop QG with graph neural networks (Pan et al., 2020) and QG for very long answers (Cao and Wang, 2021). Moreover, QG models are used to attain better interpretability in question answering such as multi-hop question decomposition (Perez et al., 2020) and question rewriting (Lee et al., 2020). As future work, we will expand our analysis to more complex scenarios and explore the connectivity with the QA task.

Ethics Statement

As the potential risk at using our QG models, it has been reported that language models inherit undesirable biases and generate toxic language (Schick et al., 2021), and one could find such text in the generated question. However, we internally checked the generated questions used for the manual evaluation, and confirmed that they did not contain toxic content.

Acknowledgements

Jose Camacho-Collados is supported by a UKRI Future Leaders Fellowship.

References

- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. [The \(un\)suitability of automatic evaluation metrics for text simplification](#). *Computational Linguistics*, 47(4):861–889.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational*

- Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, et al. 2020. Unilmv2: Pseudo-masked language models for unified language model pre-training. In *International Conference on Machine Learning*, pages 642–652. PMLR.
- Max Bartolo, Tristan Thrush, Robin Jia, Sebastian Riedel, Pontus Stenetorp, and Douwe Kiela. 2021. Improving question answering model robustness with synthetic adversarial data generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8830–8848, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Manik Bhandari, Pranav Narayan Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. Re-evaluating evaluation in text summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9347–9359, Online. Association for Computational Linguistics.
- Johannes Bjerva, Nikita Bhutani, Behzad Golshan, Wang-Chiew Tan, and Isabelle Augenstein. 2020. SubjQA: A Dataset for Subjectivity and Review Comprehension. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5480–5494, Online. Association for Computational Linguistics.
- Shuyang Cao and Lu Wang. 2021. Controllable opened question generation with a new question type ontology. In *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. Association for Computational Linguistics.
- Carrino Casimiro Pio, Costa-jussa Marta R., and Fonolosa Jose A. R. 2019. Automatic Spanish Translation of the SQuAD Dataset for Multilingual Question Answering. *arXiv e-prints*, page arXiv:1912.05200v1.
- Ying-Hong Chan and Yao-Chung Fan. 2019. A recurrent BERT-based model for question generation. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 154–162, Hong Kong, China. Association for Computational Linguistics.
- Yiran Chen, Zhenqiao Song, Xianze Wu, Danqing Wang, Jingjing Xu, Jiaye Chen, Hao Zhou, and Lei Li. 2021. Mtg: A benchmarking suite for multilingual text generation. *arXiv preprint arXiv:2108.07140*.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Danilo Croce, Alexandra Zelenanska, and Roberto Basili. 2018. Neural learning for question answering in italian. In *AI*IA 2018 – Advances in Artificial Intelligence*, pages 389–402, Cham. Springer International Publishing.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 376–380, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Martin d’Hoffschmidt, Wacim Belblidia, Quentin Heinrich, Tom Brendlé, and Maxime Vidal. 2020. FQuAD: French question answering dataset. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1193–1208, Online. Association for Computational Linguistics.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in Neural Information Processing Systems*, 32.
- Xinya Du and Claire Cardie. 2018. Harvesting paragraph-level question-answer pairs from Wikipedia. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1907–1917, Melbourne, Australia. Association for Computational Linguistics.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. Learning to ask: Neural question generation for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.
- Pavel Efimov, Andrey Chertok, Leonid Boytsov, and Pavel Braslavski. 2020. Sberquad–russian reading comprehension dataset: Description and analysis. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 3–15. Springer.

- Michael Heilman and Noah A. Smith. 2010. [Good question! statistical ranking for question generation](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617, Los Angeles, California. Association for Computational Linguistics.
- Gautier Izacard and Edouard Grave. 2020a. [Distilling knowledge from reader to retriever for question answering](#).
- Gautier Izacard and Edouard Grave. 2020b. [Leveraging passage retrieval with generative models for open domain question answering](#).
- Robin Jia, Mike Lewis, and Luke Zettlemoyer. 2021. Question answering infused pre-training of general-purpose contextualized representations. *arXiv preprint arXiv:2106.08190*.
- Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. 2019. Information maximizing visual question generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2008–2018.
- Igor Labutov, Sumit Basu, and Lucy Vanderwende. 2015. [Deep questions without deep understanding](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 889–898, Beijing, China. Association for Computational Linguistics.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Dong Bok Lee, Seanie Lee, Woo Tae Jeong, Donghwan Kim, and Sung Ju Hwang. 2020. [Generating diverse and consistent QA pairs from contexts with information-maximizing hierarchical conditional VAEs](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 208–224, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Ludovic Denoyer, and Sebastian Riedel. 2019. [Unsupervised question answering by cloze translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4896–4910, Florence, Italy. Association for Computational Linguistics.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020b. [MLQA: Evaluating cross-lingual extractive question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. [PAQ: 65 million probably-asked questions and what you can do with them](#). *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- Seungyoung Lim, Myungji Kim, and Jooyoul Lee. 2019. Korquad1.0: Korean qa dataset for machine reading comprehension. *arXiv preprint arXiv:1909.07005*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- David Lindberg, Fred Popowich, John Nesbit, and Phil Winne. 2013. [Generating natural language questions to support learning on-line](#). In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 105–114, Sofia, Bulgaria. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Luis Enrico Lopez, Diane Kathryn Cruz, Jan Christian Blaise Cruz, and Charibeth Cheng. 2020. [Transformer-based end-to-end question generation](#). *arXiv preprint arXiv:2005.01107*, 4.
- John Miller, Karl Krauth, Benjamin Recht, and Ludwig Schmidt. 2020. The effect of natural distribution shift on question answering models. In *International Conference on Machine Learning*, pages 6905–6916. PMLR.
- Ruslan Mitkov and Le An Ha. 2003. [Computer-aided generation of multiple-choice tests](#). In *Proceedings of the HLT-NAACL 03 Workshop on Building Educational Applications Using Natural Language Processing*, pages 17–22.
- Timo Möller, Julian Risch, and Malte Pietsch. 2021. [Germanquad and germandpr: Improving non-english question answering and passage retrieval](#).

- Preksha Nema and Mitesh M. Khapra. 2018. [Towards a better metric for evaluating question generation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3950–3959, Brussels, Belgium. Association for Computational Linguistics.
- Liangming Pan, Yuxi Xie, Yansong Feng, Tat-Seng Chua, and Min-Yen Kan. 2020. [Semantic graphs for generating deep questions](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1463–1475, Online. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Bhargavi Paranjape, Matthew Lamm, and Ian Tenney. 2021. [Retrieval-guided counterfactual generation for qa](#). *arXiv preprint arXiv:2110.07596*.
- Ethan Perez, Patrick Lewis, Wen-tau Yih, Kyunghyun Cho, and Douwe Kiela. 2020. [Unsupervised question decomposition for question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8864–8880, Online. Association for Computational Linguistics.
- Raul Puri, Ryan Spring, Mohammad Shoeybi, Mostofa Patwary, and Bryan Catanzaro. 2020. [Training question answering models from synthetic data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5811–5826, Online. Association for Computational Linguistics.
- Valentina Pyatkin, Paul Roit, Julian Michael, Yoav Goldberg, Reut Tsarfaty, and Ido Dagan. 2021. [Asking it all: Generating contextualized questions for any semantic role](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1429–1441, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. 2020. [ProphetNet: Predicting future n-gram for sequence-to-sequence Pre-training](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2401–2410, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21:1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- Vasile Rus, Brendan Wyse, Paul Piwek, Mihai Lintean, Svetlana Stoyanchev, and Christian Moldovan. 2010. [The first question generation shared task evaluation challenge](#). In *Proceedings of the 6th International Natural Language Generation Conference*. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. [Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in NLP](#). *Transactions of the Association for Computational Linguistics*, 9:1408–1424.
- Siamak Shakeri, Cicero Nogueira dos Santos, Henghui Zhu, Patrick Ng, Feng Nan, Zhiguo Wang, Ramesh Nallapati, and Bing Xiang. 2020. [End-to-end synthetic data generation for domain adaptation of question answering systems](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5445–5460, Online. Association for Computational Linguistics.
- ByungHoon So, Kyuhong Byun, Kyungwon Kang, and Seongjin Cho. 2022. [Jaquad: Japanese question answering dataset for machine reading comprehension](#). *arXiv preprint arXiv:2202.01764*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). *Advances in neural information processing systems*, 27.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordoni, Philip Bachman, and Kaheer Suleman. 2017. [NewsQA: A machine comprehension dataset](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 191–200, Vancouver, Canada. Association for Computational Linguistics.
- George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, et al. 2015. [An overview of the bioasq large-scale biomedical semantic indexing and question answering competition](#). *BMC bioinformatics*, 16(1):1–28.
- Shuohang Wang and Jing Jiang. 2016. [Machine comprehension using match-1stm and answer pointer](#). *arXiv preprint arXiv:1608.07905*.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Dongling Xiao, Han Zhang, Yukun Li, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. Ernie-gen: an enhanced multi-flow pre-training and fine-tuning framework for natural language generation. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3997–4003.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Zhilin Yang, Bhuwan Dhingra, Ye Yuan, Junjie Hu, William W Cohen, and Ruslan Salakhutdinov. 2017. Words or characters? fine-grained gating for reading comprehension. In *ICLR (Poster)*.
- Shiyue Zhang and Mohit Bansal. 2019. [Addressing semantic drift in question generation for semi-supervised question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2495–2509, Hong Kong, China. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. [MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 563–578, Hong Kong, China. Association for Computational Linguistics.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671. Springer.

A Parameter Optimization

A.1 Best Parameters

Model	Epoch	Learning Rate	Batch	Gradient Steps	Label Smoothing
<i>Answer-aware Model (paragraph-level)</i>					
BART _{BASE}	7	0.0001	32	8	0.15
BART _{LARGE}	4	0.00005	32	4	0.15
T5 _{SMALL}	9	0.0001	64	1	0.15
T5 _{BASE}	5	0.0001	16	4	0.15
T5 _{LARGE}	6	0.00005	16	4	0.15
<i>Answer-aware Model (sentence-level)</i>					
BART _{BASE}	3	0.0001	64	2	0.15
BART _{LARGE}	8	0.00005	32	16	0.15
T5 _{SMALL}	8	0.0001	64	1	0.15
T5 _{BASE}	8	0.0001	64	1	0.15
T5 _{LARGE}	6	0.00005	16	4	0.15
<i>Answer-free Model</i>					
BART _{BASE}	4	0.0001	32	8	0.15
BART _{LARGE}	4	0.00005	32	4	0.15
T5 _{SMALL}	7	0.0001	64	4	0.15
T5 _{BASE}	8	0.0001	16	4	0.15
T5 _{LARGE}	7	0.00005	16	4	0.15

Table 7: The best parameter to fine-tune each model on SQuAD we found through the parameter optimization.

Table 7 shows the best configuration to fine-tune each model that we obtain through the parameter optimization process. To fine-tune T5 model, we use the task prefix `generate question:` at the beginning of the input text.

A.2 Fine-tuning without Optimization

Model	B4	R-L	MTR	BS	MS
BART _{BASE}	-0.28	-0.17	-0.07	-0.01	0.00
BART _{LARGE}	-2.22	-1.65	-1.16	-0.06	-0.57
T5 _{SMALL}	-1.73	-1.89	-1.16	-0.28	-0.82
T5 _{BASE}	-0.72	-0.58	-0.39	-0.10	-0.28
T5 _{LARGE}	-0.18	-0.15	0.00	-0.07	-0.09

Table 8: Decrease in automatic metrics of our QG models without parameter optimization.

Table 8 shows the decrease in each metric for SQuAD if the model is fine-tuned without parameter optimization.¹² We observe decent drops in performance. T5_{SMALL} and BART_{LARGE} lose around 2 points in BLEU4 and ROUGE_L. According to these results, we infer that T5 and BART were worse than more recent LMs (ProphetNet, UniLM, or ERNIE-GEN) in QG just because they were under-fitted to the task due to sub-optimal fine-tuning parameters, rather than they being inferior

¹²We follow the hyperparameter used to fine-tune ERNIE-GEN on SQuAD QG in the original paper.

to those recent LMs in terms of learning the QG task.

B Manual Evaluation

B.1 Sample Outputs

Table 9 presents a few examples of our model predictions with the scores made by the annotators, where the samples are chosen from the high-answerability and low-answerability groups of T5_{LARGE}.

B.2 Spearman’s Correlation

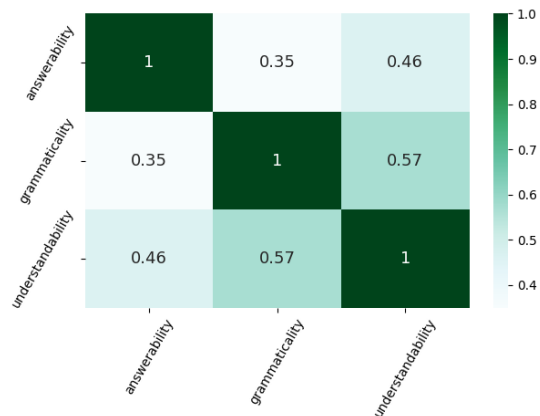


Figure 5: Spearman’s rank correlation within manual evaluation criteria.

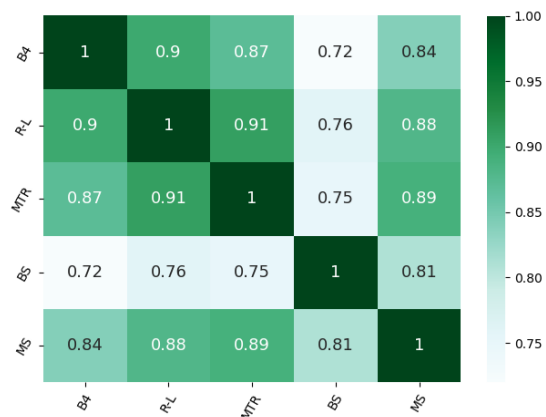


Figure 6: Spearman’s rank correlation within automatic evaluation metrics among the 500 samples we used in SQuAD manual annotation.

Figure 5 and Figure 6 show Spearman’s rank correlation across automatic metrics and manual evaluation criteria among the questions we generate over SQuAD test set for the manual annotation. The p-values of all those correlation are less than 0.05 so they are statistically significant.

Paragraph	Question	Model	Prediction	Answerability	Grammaticality	Understandability
<p>The British Library notes that "Chopin's works have been recorded by all the great pianists of the recording era." The earliest recording was an 1895 performance by Paul Pabst of the Nocturne in E major Op. 62 No. 2. The British Library site makes available a number of historic recordings, including some by Alfred Cortot, Ignaz Friedman, Vladimir Horowitz, Benno Moiseiwitsch, Paderewski, Arthur Rubinstein, Xaver Scharwenka and many others. A select discography of recordings of Chopin works by pianists representing the various pedagogic traditions stemming from Chopin is given by Methuen-Campbell in his work tracing the lineage and character of those traditions.</p>	<p>When did Pabst record his Chopin performance?</p>	NQG	What year was the earliest recording of the earliest recording?	1	2.6	1.6
		BART _{LARGE}	When was the earliest recording of Chopin's Nocturne in E major?	2.8	3	3
		T5 _{SMALL}	When did Paul Pabst perform the Nocturne in E major Op. 62 No. 2?	2.8	3	3
		T5 _{LARGE}	When was the earliest recording of Chopin's Nocturne in E major Op. 62 No. 2?	3	3	3
		-sent-level	When was the earliest recording of Beethoven's Nocturne in E major?	1	3	3
		-answer-free	In what year did Paul Pabst record Chopin's Nocturne in E major Op. 62 No. 2?	3	3	3
<p>Baird Television Ltd. made Britain's first television broadcast, on 30 September 1929 from its studio in Long Acre, London, via the BBC's London transmitter, using the electromechanical system pioneered by John Logie Baird. This system used a vertically-scanned image of 30 lines – just enough resolution for a close-up of one person, and with a bandwidth low enough to use existing radio transmitters. Simultaneous transmission of sound and picture was achieved on 30 March 1930, by using the BBC's new twin transmitter at Brookmans Park. By late 1930, 30 minutes of morning programmes were broadcast Monday to Friday, and 30 minutes at midnight on Tuesdays and Fridays, after BBC radio went off the air. Baird broadcasts via the BBC continued until June 1932.</p>	<p>How many lines made up the picture on Britain's earliest TV broadcasts?</p>	NQG	When did the UK's first television broadcast?	1	2.6	2.2
		BART _{LARGE}	How many lines did Baird Television Ltd.'s first television broadcast use?	3	3	2.6
		T5 _{SMALL}	On what date did Baird Television make Britain's first television broadcast?	1	3	3
		T5 _{LARGE}	On what date did Baird Television Ltd. make Britain's first television broadcast?	1	3	3
		-sent-level	When was Britain's first television broadcast?	1	3	3
		-answer-free	When did Baird Television Ltd. make Britain's first television broadcast?	1	3	2.8

Table 9: Examples of the system outputs along with their scores from the manual evaluation. The sentence and answer are highlighted by boldface and underline in the paragraph.

Data	F1	Exact Match
BART _{BASE}	70.10	58.46
BART _{LARGE}	70.40	58.60
T5 _{SMALL}	68.90	56.96
T5 _{BASE}	70.33	58.14
T5 _{LARGE}	70.86	59.04

Table 10: Unsupervised QA-based evaluation results of our answer-aware QG models (paragraph-level). All results are the performance on the validation set of original SQuAD by the model trained on the synthetic data generated by each QG model.

B.3 William test

In § 6.2, we run correlation analysis and here we report the result of the William test to check the significance of the increase in the correlation across metrics in Figure 7, showing that the increase is statistically significant as well.

B.4 Guidelines

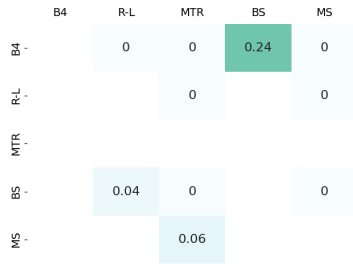
Figure 8 shows an example of user interface we implemented for our manual evaluation and the guideline we present to the annotators is attached to the end of the paper.

C Unsupervised QA-based Evaluation

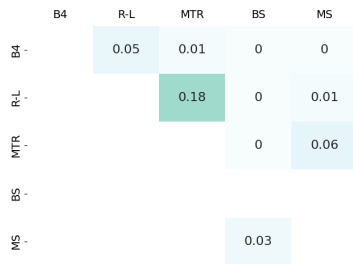
As a proxy for *answerability*, we run an unsupervised QA-based evaluation (Zhang and Bansal, 2019), which trains a QA model on synthetic data

generated by the target QG model and evaluates the QA model on a human annotated test set. As an alternative to the traditional metrics in QG, Q-metric (Nema and Khapra, 2018) shows high agreement in terms of the *answerability*, but we prefer to employ QA-based evaluation (Zhang and Bansal, 2019), since it is more closely tied to downstream applications, while Q-metric relies on some heuristics such as the number of named-entity/pre-defined question types. This evaluates the QG model's capability to generate high quality questions: higher accuracy of the QA model indicates a better QG model. The synthetic data is usually generated over the paragraph and answer (PA) pairs collected by Du and Cardie (2018). Zhang and Bansal (2019) used a small subset of the PA pairs, since they contain 12x larger instances than the SQuAD training set. Since this introduces an artifact of the subset choice, we decided to train QA models on the entire PA pairs set with the generated questions. Also, we train QA models solely on the synthetic data, which differs from work in semi-supervised QA where the QA model is trained on a concatenation of the synthetic data and the original SQuAD training set (Lee et al., 2020).

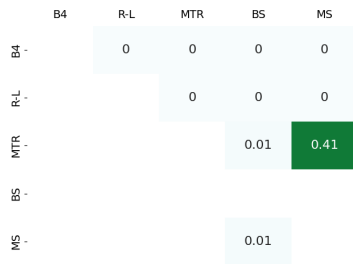
The synthetic QA data is created by generating a question for each of the one million PA pairs (Du and Cardie, 2018) with the target QG model.



(a) answerability



(b) grammaticality



(c) understandability

Figure 7: Williams test on the difference in the correlation reported in Figure 3. The difference of correlation is significant if the value is less than 0.005.

We then fine-tune BERT (Devlin et al., 2019)¹³ on the synthetic QA data with the default configuration used in the HuggingFace’s tutorial to fine-tune BERT on QA.¹⁴ We report F1 score and the exact match on the SQuAD validation set, following Zhang and Bansal (2019).¹⁵

The results of our unsupervised QA-based evaluation in Table 10 indicate that the QA model accuracy correlates with the size of QG model that generated the synthetic data, as in T5_{LARGE} realizes the best QA model in both of F1 and the exact match, which is as good as the supervised non-language model based QA models (Wang and Jiang, 2016;

¹³We use bert-base-cased from HuggingFace.

¹⁴<https://github.com/huggingface/transformers/tree/master/examples/pytorch/question-answering>

¹⁵We will release the synthetic data we made on Huggingface Dataset <https://huggingface.co/datasets>.

Yang et al., 2017). Also, the small models such as T5_{SMALL} and BART_{BASE} produce QA models with a small decrease in performance, which exhibits the efficiency of our models, similarly to our results with automatic metrics.

D Additional Analysis

D.1 Zero-shot Multilingual Transfer

Data	B4	R-L	MTR	BS	MS
SQuAD	21.65	48.95	23.83	90.01	62.75
Ru	0.00	0.99	1.78	70.89	49.10
Ja	0.00	6.08	0.51	66.08	46.53
It	0.54	5.01	5.89	72.60	50.23
Ko	0.00	0.06	0.73	66.34	45.86
Es	0.59	5.21	6.02	74.94	50.62
De	0.00	1.56	4.81	73.53	50.37
Fr	1.71	15.84	8.24	72.91	50.96

Table 11: Zero-shot result of mT5 fine-tuned on SQuAD except for the first row, which shows fine-tuning result of SQuAD.

We fine-tune multilingual language model on each of multilingual QG dataset in § 5.2, and here we explore the zero-shot multilingual transfer by evaluating English fine-tuned QG model in other languages. Table 11 shows the zero-shot transfer result where we fine-tune mT5_{SMALL} on SQuAD and evaluate it on the test set of multilingual QG dataset. Compared with Table 3, the performance is largely decreased, indicating the difficulty of zero-shot multilingual transfer in QG.

D.2 Zero-shot Domain Transfer

Figure 9 shows the comparison of zero-shot QG transfer in SQuADShifts and SubjQA dataset with T5_{LARGE}.

Question Evaluation

In this project, we aim to study the quality of questions generated by automatic systems. You will be given the following 3 pieces of information to complete the evaluation.

- **[Passage]:** Passage consisting of multiple sentences with the information required editable to answer the question. The sentence that should contain the answer to the question is boldfaced.
- **[Answer]:** Answer to the question. This is usually an entity that appears in the passage.
- **[Question A - F]:** Questions written by our systems based on the passage and the answer.

We ask you to evaluate each of the 6 questions based on the following 4 criteria with a 3-point scale:

- **Grammaticality** is the grammatical correctness of the question (do not refer the answer or the passage).
 - 3: correct
 - 2: minor errors
 - 1: major errors
- **Understandability** is how understandable the question is given the passage.
 - 3: easy to understand
 - 2: complicated yet understandable
 - 1: too complex to understand
- **Correctness** is whether the answer to the question matches given the passage.
 - 3: the question is asking the answer exactly
 - 2: the question might be asking the answer but no clear evidence found in the passage
 - 1: the question is not asking the answer

BEFORE START: Each criterion is explained further in our guideline [here](#), so **please make sure to read it carefully before start**. There are a few questions in each HIT, specifically designed to check the quality of the HIT. We will manually check the evaluation made on them, and reject the HIT if the quality is too bad.

[Passage]: *Environmental sustainability has become a mainstream issue, with profound effect on the architectural profession. Many developers, those who support the financing of buildings, have become educated to encourage the facilitation of environmentally sustainable design, rather than solutions based primarily on immediate cost. Major examples of this can be found in Passive solar building design, greener roof designs, biodegradable materials, and more attention to a structure's energy usage. **This major shift in architecture has also changed architecture schools to focus more on the environment.** Sustainability in architecture was pioneered by Frank Lloyd Wright, in the 1960s by Buckminster Fuller and in the 1970s by architects such as Ian McHarg and Sim Van der Ryn in the US and Brenda and Robert Vale in the UK and New Zealand. There has been an acceleration in the number of buildings which seek to meet green building sustainable design principles. Sustainable practices that were at the core of vernacular architecture increasingly provide inspiration for environmentally and socially sustainable contemporary techniques. The U.S. Green Building Council's LEED (Leadership in Energy and Environmental Design) rating system has been instrumental in this.*

[Answer]: *the environment*

[Question A]: *What has the major shift in architecture schools to focus more on?*

- Grammaticality
 - 1
- Understandability
 - 1
- Correctness
 - 1

Figure 8: An example of the interface used in our manual evaluation.

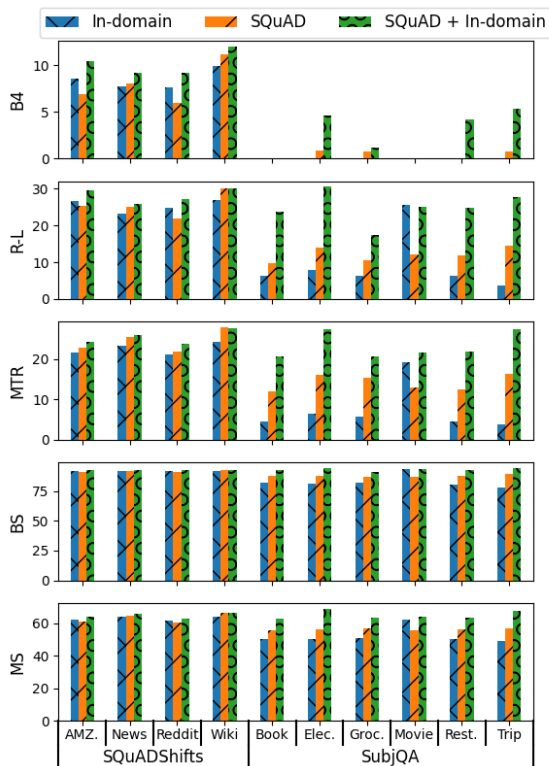


Figure 9: Metric comparison for $T5_{LARGE}$ across in-domain fine-tuning, zero-shot transfer of SQuAD fine-tuned model, and in-domain fine-tuning from SQuAD model.

Question Evaluation Guideline

In this project, we aim to study the quality of questions generated by automatic models. You will be given the following 4 pieces of information to complete the evaluation.

- **Passage:** Passage consisting of multiple sentences with the information required to answer the question.
- **Answer:** Answer to the question. This is usually an entity that appears in the passage.
- **Question:** Question written by our model based on the passage and the answer.

Goal

We ask you to evaluate *questions* based on the following 4 criteria: (1) Grammaticality, (2) Understandability, (3) Correctness, and (4) Question Difficulty.

(1) Grammaticality

You will score a question in terms of its grammatical correctness with a 3-point scale.

- 3: The question is grammatically correct.
- 2: The question has some minor errors/typos but you can still understand it.
- 1: The question contains many errors and you can not understand it.

You should **only rely on the question** and do not refer to any other information such as the passage and the answer.

(2) Understandability

You will score how understandable a question is with a 3-point scale.

- 3: The question is easy to read and understand what you are being asked.
- 2: The question is somewhat complicated to understand, but you can get an idea of what the question is asking.
- 1: The question is too complex and you can not understand what should be the answer to the question.

Understandability could correlate with grammaticality, but **a question without any grammatical mistakes can have low understandability** due to an over complex structure. Likewise, a question can be easy to understand even with a few grammatical mistakes. You can refer the passage if it's needed.

(3) Correctness

We generate a question in a way that its answer should be the answer. Here, you will evaluate whether the answer to The question matches the answer, given the passage.

- 3: The answer to the question is exactly the given answer.
- 2: The answer to the question might be the given answer but no clear evidence can be found in the passage, or the question is too vague. Or the question is not relevant to the passage.
- 1: The answer to the question is not the given answer.

To better understand when you judge a question as 2, let's look at the following example:

- Passage: Max was raised in LA ...
- Answer: LA
- Question: Where was Max born?

While the answer "LA" could be an answer to the question, it is not entirely accurate **since the passage does not explicitly state that Max was born in LA**. A more correct question could be "Where was Max raised?". In these situations, you could score the question with a 2.

The answer can sometimes be partial but it is fine. The answer in the following example should be 'June 6th 1992' instead of 'June' but you should score it with a 3 as the question still makes sense with the answer June.

- Passage: Max was born on June 6th 1992, and ...
- Answer: June
- Question: When was Max born?

In some cases, the question matches the answer while it is completely irrelevant to the passage. For example,

- Passage: China spans five geographical time zones and borders 14 different countries, the second most of any country in the world after Russia.
- Answer: China
- Question: What is the name of the world's most populous country?

Although the question matches the answer, it is based on common knowledge rather than any evidence found in the passage, so this question should be scored as 2.

In addition, if the Understandability of the question is 1, you should mark its Correctness as 1.