

LIQUID: A Framework for List Question Answering Dataset Generation

Seongyun Lee^{*1}, Hyunjae Kim^{*1}, Jaewoo Kang^{1,2}

¹Korea University,

²AIGEN Sciences

{sy-lee, hyunjae-kim, kangj}@korea.ac.kr

Abstract

Question answering (QA) models often rely on large-scale training datasets, which necessitates the development of a data generation framework to reduce the cost of manual annotations. Although several recent studies have aimed to generate synthetic questions with single-span answers, no study has been conducted on the creation of list questions with multiple, non-contiguous spans as answers. To address this gap, we propose LIQUID, an automated framework for generating list QA datasets from unlabeled corpora. We first convert a passage from Wikipedia or PubMed into a summary and extract named entities from the summarized text as candidate answers. This allows us to select answers that are semantically correlated in context and is, therefore, suitable for constructing list questions. We then create questions using an off-the-shelf question generator with the extracted entities and original passage. Finally, iterative filtering and answer expansion are performed to ensure the accuracy and completeness of the answers. Using our synthetic data, we significantly improve the performance of the previous best list QA models by exact-match F1 scores of 5.0 on MultiSpanQA, 1.9 on Quoref, and 2.8 averaged across three BioASQ benchmarks.

1 Introduction

Extractive question answering (QA) refers to the task of finding answers to questions in the provided text. Because building a QA system often requires a vast number of human-annotated training examples, recent studies have attempted to reduce annotation costs by generating synthetic datasets from unlabeled corpora (Yang et al. 2017; Dhingra, Danish, and Rajagopal 2018; Alberti et al. 2019; Lyu et al. 2021). However, these studies have focused only on generating questions with single-span answers and failed to cover list questions that require multiple spans as answers (Voorhees et al. 2001; Tsatsaronis et al. 2015). Although list questions constitute a large portion of the questions asked in practice (Yoon et al. 2022), the automatic generation of list QA datasets has not been sufficiently studied.

QA dataset generation frameworks typically consist of answer extraction, question generation, and filtering models that are trained with numerous human-labeled data (Alberti et al. 2019; Puri et al. 2020; Shakeri et al. 2020; Lewis

^{*}These authors contributed equally.

Passage: Katherine Saltzberg is an American actress, singer, and comic. She is best known for starring as the showbiz-talented 16-year-old daughter of Brian Dennehy's character in the ABC sitcom, Star of the Family. In 2009, Saltzberg wrote and performed the one woman show, Los Angelyne, ..., as she recounted how her life and home were invaded by Los Angeles icon Angelyne. ...

Summary: Katherine Saltzberg is an American actress, singer, and comic. She is best known for starring as the showbiz-talented 16-year-old daughter of Brian Dennehy's character.

Question Generation

Entities from *passage* + *passage* → Question A:
 Katherine Saltzberg, Brian Dennehy, Saltzberg, Angelyne
 What are the names of the actors?

Entities from *summary* + *passage* → Question B:
 Katherine Saltzberg, Brian Dennehy
 Who are the two actors that starred in Star of the Family?

Figure 1: Questions generated conditionally for the same passage with different sets of entities. When all the named entities in the passage are used, question A about a trivial fact covering all the entities is generated, which may not be effective for training list QA models. In contrast, using entities from a summary creates a more specific and clearer question, question B, because the entities within the summary are usually related by a common topic and fact.

et al. 2021). Unfortunately, this supervised approach is not applicable to list QA because existing large QA datasets contain only a few or no list-type questions (Rajpurkar et al. 2016; Trischler et al. 2017; Joshi et al. 2017; Rajpurkar, Jia, and Liang 2018; Yang et al. 2018; Kwiatkowski et al. 2019). Moreover, the automatic generation of list QA datasets presents unique challenges that existing frameworks cannot address. First, semantically correlated candidate answers need to be carefully selected because unrelated or excessive candidate answers can result in broad and trivial questions that may not be useful for obtaining list QA models, as depicted in Figure 1. In addition, all answers should be accurate and complete, that is, the answer set should not contain incorrect answers or omit correct answers.

In this paper, we present a new framework, **LIQUID**,

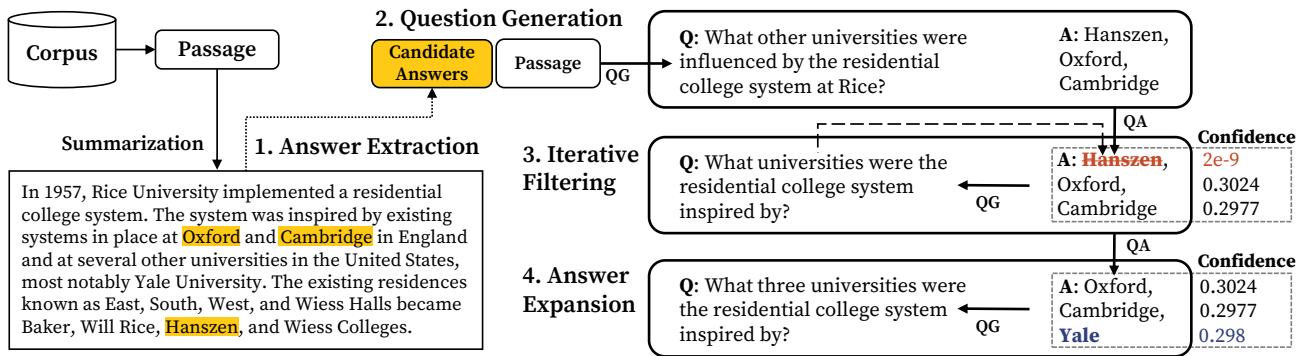


Figure 2: Overview of LIQUID. (1) Answer extraction: the named entities belonging to the same entity type (e.g., organization type) in a summary are extracted by an NER model and used as candidate answers. (2) Question generation: the candidate answers and the original passage are fed into a QG model to generate list questions. (3) Iterative filtering: incorrect answers (e.g., Hanszen) are iteratively filtered based on the confidence score assigned by a QA model. (4) Answer expansion: correct but omitted answers (e.g., Yale) are identified by the QA model. A threshold value of 0.1 was used in this example.

to automatically create **list question answering datasets**. Specifically, we first collect passages from unlabeled corpora, such as Wikipedia and PubMed. Subsequently, a summarization model is used to convert these passages into summaries, and named entity recognition (NER) models extract answers from these summaries. As entities in the same summary are likely to be semantically correlated to a common topic or fact (Lu et al. 2022), they can be used as suitable candidate answers for list-type questions. For instance, Figure 1 illustrates that the entities in the summary (i.e., “Katherine Saltzberg” and “Brian Dennehy”) have a common characteristic in that they appear in the sitcom “Star of the Family,” which enables the production of a specific and clear question. To generate questions, we input the extracted entities and passage into a question generation (QG) model trained on SQuAD (Rajpurkar et al. 2016). After the QG process, we improve the quality of the question-answer pairs using a QA model trained on SQuAD. We perform iterative filtering to ensure accuracy, wherein the QA model is used to eliminate answers with confidence scores lower than a threshold, and the QG model is used to re-generate questions based on the passage and new answer set. These processes are iterated until the answer set remains unchanged. For completeness, we perform answer expansion to determine additional spans, which the QA model assigns confidence scores that are higher than the lowest score in the answer set. The entire process of LIQUID allows us to obtain large-scale list QA datasets without relying on hand-labeled list QA data to train each component.

We used five datasets comprising MultiSpanQA (Li et al. 2022) and Quoref (Dasigi et al. 2019) for the general domain, and the BioASQ 7b (Nentidis et al. 2019), 8b (Nentidis et al. 2020), and 9b (Nentidis et al. 2021) datasets for the biomedical domain. When the models were trained using our synthetic data, the exact-match F1 scores improved by 5.0 on MultiSpanQA, 1.9 on Quoref, and 2.8 on the three BioASQ datasets, compared with the scores obtained using only human-labeled data. We conducted extensive ablation studies to confirm the effectiveness of our proposed meth-

ods. In addition, we validated the quality of the generated data and discussed their limitations.

In summary, we made the following contributions:

- To the best of our knowledge, this is the first study to introduce a framework for list QA dataset generation. We addressed the unique challenges of creating list QA datasets by combining current large-scale models with several methods such as summarization-based answer extraction, iterative filtering, and answer expansion.
- We significantly improved the performance of the previous best models by an F1 score of 1.9 to 5.0 on five datasets in the general and biomedical domains.
- Our code and data have been made publicly available to facilitate further research and real-world applications.¹

2 List QA Dataset Generation

Our goal is to automatically generate list QA data $\tilde{\mathcal{D}}$ from an unlabeled corpus \mathcal{C} to supplement human-labeled data \mathcal{D} . Notably, our framework consists of (1) answer extraction, (2) question generation, (3) iterative filtering, and (4) answer expansion. The initial data are generated in stages 1 and 2, whereas data refinement is performed in stages 3 and 4. Figure 2 presents the overall process, and Algorithm 1 (Appendix) details the process using a formal notation.

Answer Extraction

Let c be a passage from corpus \mathcal{C} . We first summarize c into \bar{c} and extracted named entities from \bar{c} . Entities of the same types are then used as candidate answers. Details pertaining to this process are described in the following sections.

Summarization One of the most important considerations when selecting candidate answers for list QA is that the answers must be semantically correlated in the given context. As noted in Section 1, unrelated candidate answers lead to trivial list questions, because the QG model attempts to construct questions that encompass all given candidate answers.

¹<https://github.com/dmis-lab/LIQUID>

However, it is difficult to select appropriate candidate answers from all possible spans in a passage. Instead, we use summarized text (i.e., \bar{c}) because a summary is a short snippet that conveys the relevant topics/facts in one or multiple documents; therefore, similar phrases (e.g., named entities of the same type) in the same summary are likely to be semantically related to one another and appropriate for list-type questions. We used the $\text{BART}_{\text{base}}$ model (Lewis et al. 2020) trained on the CNN/Daily Mail dataset (Nallapati et al. 2016) as the summarization model.

NER The answers to list questions usually comprise named entities (see Section 4). In addition, in most cases, the answers to a given question have the same entity type. Thus, we extract the named entities from a given text and use them as candidate answers. We regard entities of the same type as belonging to the same group of answers. Formally, we obtain L sets of candidate answers $\mathbf{A}_1, \dots, \mathbf{A}_L$ for each summary \bar{c} , where L denotes the number of predefined entity types and $\mathbf{A}_l = \{a_{l_1}, \dots, a_{l_M}\}$ denotes the l -th set of candidate answers consisting of l_M entities. We omit subscript l for simplicity in the following sections. We used the spaCy² NER tagger and BERN2 (Sung et al. 2022) for the general and biomedical domains, respectively.

Question Generation

Previous studies attempted to generate questions by training sequence-to-sequence models (Radford et al. 2019; Lewis et al. 2020; Raffel et al. 2020). Specifically, in these studies, the QG model considered a single candidate answer a with passage c as an input, which was represented as “answer: a context: c .” The model was optimized to generate the corresponding question q , where the triplet $\langle c, q, a \rangle$ was obtained from large-scale QA datasets. We adopt this approach with a simple modification to the input format, wherein we concatenate all the candidate answers (i.e., named entities) with commas and prepend them to passage c as follows: “answer: a_1, \dots, a_M context: c .” The input is then fed into a T5_{base} model (Raffel et al. 2020) trained on a single-span QA dataset, SQuAD. Interestingly, the model generates appropriate questions, despite not being trained to generate list questions for multiple given answers (see Section 4). After the QG process, we obtain the initial QA instance $\langle c, q, \mathbf{A} \rangle$, where $\mathbf{A} = \{a_1, \dots, a_M\}$ is the set of M answers.

Iterative Filtering

Filtering Because the initial answer set \mathbf{A} can contain incorrect answers to question q , we verify the answer set for better accuracy. Given passage c and question q , a single-span QA model is used to obtain the confidence scores for all the candidate answers a_1, \dots, a_M . Answers with confidence scores lower than the threshold τ are regarded as incorrect and filtered out to yield a refined answer set $\mathbf{A}' = \{a'_1, \dots, a'_{M'}\}$ ($M' \leq M$). The triplet $\langle c, q, \mathbf{A}' \rangle$ is not used if zero or one answer remains after filtering (i.e., $M' \leq 1$); otherwise, the instance is passed to the next stage. For the QA model, we used $\text{RoBERTa}_{\text{base}}$ (Liu et al. 2019)

²<https://spacy.io/>

Dataset	Train	Valid	Test
<i>General domain</i>			
MultiSpanQA (Li et al. 2022)	5,230	653	653
Quoref (Dasigi et al. 2019)	1,766	221	221
<i>Biomedical domain</i>			
BioASQ 7b (Nentidis et al. 2019)	556	88	88
BioASQ 8b (Nentidis et al. 2020)	644	75	75
BioASQ 9b (Nentidis et al. 2021)	719	94	94

Table 1: Number of questions in list QA benchmark datasets. “Train,” “valid,” and “test” indicate the training, validation, and test sets, respectively.

or $\text{BioBERT}_{\text{base}}$ (Lee et al. 2020) with linear prediction layers for the general and biomedical domains, respectively. Both models were trained using SQuAD.

Question re-generation The initial question q may not perfectly align with the answers obtained after filtering. Therefore, we re-generate question q' based on answer set \mathbf{A}' and passage c in the same manner as that described in Section 2. Filtering is then performed, and the process is repeated until the current answer set is the same as the previous one or the maximum number of iterations T is reached.

Specifying answer positions The start and end positions of the answers are required to train QA models. Because we extract the answers from summary \bar{c} and use passage c as the evidence text, the correct positions for the answers in the original passage need to be identified. We address this problem by using the span with the highest confidence score of the QA model as the answer position for an answer string.

Answer Expansion

The initial set \mathbf{A} is often incomplete, primarily owing to false negatives generated by the NER model. Although the accuracy of the generated data is improved by excluding incorrect answers during the filtering process, this incompleteness cannot be addressed through filtering. We address this issue by identifying additional answer spans with confidence scores higher than the lowest confidence in the filtered set \mathbf{A}' by using the same QA model as that used for filtering. A question q'' based on the expanded set \mathbf{A}'' is then generated, and triplet $\langle c, q'', \mathbf{A}'' \rangle$ is used as the final instance if the QA model does not filter any answers in \mathbf{A}'' for question q'' ; otherwise, triplet $\langle c, q', \mathbf{A}'' \rangle$ is used.

3 Experiments

Datasets

We used (1) MultiSpanQA (Li et al. 2022), which consists of English Wikipedia passages with list questions and answers. The passages and questions were selected from the Natural Questions dataset (Kwiatkowski et al. 2019), and the answers were re-annotated by humans to improve their quality. Note that the model evaluation on the hold-out test set

Model	MultiSpanQA		Quoref	
	Exact F1 (P/R)	Partial F1 (P/R)	Exact F1 (P/R)	Partial F1 (P/R)
<i>Baselines: labeled only (\mathcal{D})</i>				
BERT _{base} + S*	14.4 (16.2/13.0)	67.6 (60.3/76.8)	-	-
BERT _{base} + T*	56.5 (52.5/61.1)	75.2 (75.9/74.5)	-	-
BERT _{base} + T (multi-task)*	59.3 (58.1/60.5)	76.3 (79.6/73.2)	-	-
RoBERTa _{base} + S	10.5 (14.4/8.3)	63.0 (60.0/66.3)	55.4 (65.2/48.0)	69.0 (76.7/62.6)
RoBERTa _{base} + T	62.9 (63.0/62.9)	78.0 (82.5/73.9)	81.2 (73.8/90.1)	85.7 (80.1/92.2)
RoBERTa _{large} + T	66.4 (62.3/71.2)	82.6 (82.1/83.0)	84.2 (76.1/94.2)	88.8 (82.6/96.0)
CorefRoBERTa _{large} + T	64.0 (56.5/73.8)	81.7 (77.7/86.0)	86.5 (81.3/92.4)	89.7 (86.1/93.7)
<i>Our models: synthetic & labeled ($\tilde{\mathcal{D}} \rightarrow \mathcal{D}$)</i>				
RoBERTa _{base} + S	19.4 (19.7/19.0)	71.0 (62.9/81.4)	60.7 (63.8/57.9)	74.3 (77.4/71.3)
RoBERTa _{base} + T	67.4 (65.7/69.2)	81.2 (80.9/81.5)	85.7 (82.3/89.3)	89.1 (86.5/91.8)
RoBERTa _{large} + T	71.4 (75.0/68.2)	80.9 (85.3/77.0)	86.7 (85.8/87.6)	90.2 (89.4/91.1)
CorefRoBERTa _{large} + T	65.8 (64.0/67.8)	80.2 (79.8/80.5)	88.4 (84.8/92.2)	91.7 (89.1/94.4)

Table 2: Exact-match and partial-match F1 scores (precision/recall) of QA models on the test sets of MultiSpanQA (Li et al. 2022) and Quoref (Dasigi et al. 2019). ‘‘S’’ and ‘‘T’’ indicate single-span extractive and sequence tagging models, respectively. * indicates that the model was implemented by Li et al. (2022).

is only available on the official leaderboard.³ We also used another Wikipedia-based dataset, (2) Quoref (Dasigi et al. 2019), which was originally designed to assess the ability to answer questions that require co-reference reasoning. The dataset contains some list questions with multiple answers. We used the original validation set as the test set and split the original training set into the training and validation sets. For the biomedical domain, we used three datasets provided in the recent (3) BioASQ challenges (Tsatsaronis et al. 2015) comprising BioASQ 7b, 8b, and 9b (Nentidis et al. 2019, 2020, 2021). These datasets comprise evidence texts from biomedical literature with manual annotations by experts. We sampled data from the training set to form a validation set. Table 1 presents the dataset statistics.

List QA Models

Because LIQUID is a model-agnostic framework, any of list QA models can be used. We used two types of models: single-span extractive and sequence tagging models (see the paragraphs below). For the text encoder, we used RoBERTa (base and large) and CorefRoBERTa_{large} (Ye et al. 2020) for the general domain and BioBERT (base and large) for the biomedical domain. In addition, for MultiSpanQA, we included three BERT-based baseline models (Devlin et al. 2019) used in the study of Li et al. (2022). Among them, the ‘‘multi-task’’ model is the previous best model on MultiSpanQA, which is trained with additional objective functions such as span number prediction and structure prediction. In the fine-tuning stage, we selected the best model checkpoints based on exact-match F1 scores for the validation set at every epoch. The maximum epochs were set to 20 and 50 for the single-span extractive and sequence tagging models, respectively. The detailed hyperparameter settings are given in Appendix B.

³<https://multi-span.github.io>

Single-span extractive model A conventional approach for performing list QA involves the use of a standard extractive QA model with an absolute threshold, wherein all spans with confidence scores exceeding a threshold are used as the predicted answers, and the threshold is a hyperparameter.

Sequence tagging model Single-span extractive models are typically unsuitable for list QA. In recent studies, the list QA problem has been formulated as a sequence labeling problem, for which models are required to predict the beginning, inside, and outside (i.e., BIO) labels for each token (Segal et al. 2020; Yoon et al. 2022).

Metrics

Following Li et al. (2022), we used strict and relaxed evaluation methods, (1) exact and (2) partial matches, respectively. Both methods are based on micro-averaged precision (P), recall (R), and F1 score (F1),⁴ which are computed as follows:

$$P = \frac{\sum_{n=1}^N \sum_{\hat{a} \in \hat{\mathbf{A}}_n} f(\hat{a}, \mathbf{A}_n^*)}{\sum_{n=1}^N |\hat{\mathbf{A}}_n|},$$

$$R = \frac{\sum_{n=1}^N \sum_{a^* \in \mathbf{A}_n^*} f(a^*, \hat{\mathbf{A}}_n)}{\sum_{n=1}^N |\mathbf{A}_n^*|}, \quad F1 = \frac{2 \cdot P \cdot R}{(P + R)},$$

where N is the number of questions, \mathbf{A}_n^* is the set of gold answers, $\hat{\mathbf{A}}_n$ is the set of predictions for the n -th question, and $|\cdot|$ is the number of elements in the set. The exact-match and partial-match evaluation methods differ in the scoring function f , as detailed in the following paragraphs.

Exact match Because the strings of model prediction should ideally be identical to those of the gold answer, the scoring function f is defined as $f(x, \mathbf{Y}) := I_{\mathbf{Y}}(x)$, where

⁴We use the terms ‘‘exact-match F1 score’’ and ‘‘F1 score’’ interchangeably.

Model	BioASQ 7b		BioASQ 8b		BioASQ 9b	
	Exact F1 (P/R)	Partial F1 (P/R)	Exact F1 (P/R)	Partial F1 (P/R)	Exact F1 (P/R)	Partial F1 (P/R)
<i>Baselines: labeled only (\mathcal{D})</i>						
BioBERT _{base} + S	42.1 (55.9/33.8)	60.2 (82.3/47.5)	34.4 (44.9/27.9)	53.5 (40.2/79.9)	56.1 (46.2/71.3)	73.8 (70.3/77.7)
BioBERT _{base} + T	46.1 (39.7/55.1)	70.5 (68.5/72.6)	41.8 (33.5/55.5)	67.6 (64.0/71.5)	66.7 (60.1/74.9)	80.6 (76.4/85.2)
BioBERT _{large} + T	49.5 (40.5/63.6)	74.6 (70.7/78.9)	45.0 (34.7/64.0)	72.2 (65.8/80.0)	68.2 (60.9/77.5)	81.4 (76.3/87.2)
<i>Our models: synthetic & labeled ($\tilde{\mathcal{D}} \rightarrow \mathcal{D}$)</i>						
BioBERT _{base} + S	51.8 (49.0/55.0)	70.2 (69.7/70.7)	44.2 (41.4/47.5)	65.2 (65.4/65.0)	64.0 (58.0/71.4)	76.6 (72.6/81.1)
BioBERT _{base} + T	49.0 (41.0/61.0)	73.1 (70.4/76.0)	44.2 (36.6/55.8)	69.4 (67.3/71.7)	71.5 (67.0/76.6)	83.2 (80.0/86.7)
BioBERT _{large} + T	52.3 (44.5/63.5)	74.9 (71.9/78.1)	46.5 (38.5/58.8)	72.3 (68.9/76.1)	72.2 (67.3/77.8)	83.4 (80.4/86.7)

Table 3: Exact-match and partial-match F1 scores (precision/recall) of QA models on the test sets of the BioASQ 7b, 8b, and 9b datasets (Nentidis et al. 2019, 2020, 2021). S: single-span extractive model. T: sequence tagging model.

$I_{\mathbf{Y}}$ is an indicator function that returns one if $x \in \mathbf{Y}$ and zero; otherwise, x is a span, and \mathbf{Y} is a set of spans.

Partial match This evaluation considers the overlap between the strings of gold answers and predictions. The scoring function f is defined as $f(x, \mathbf{Y}) := \max_{y \in \mathbf{Y}} (g(x, y) / \text{len}(x))$, where $g(x, y)$ is the length of the longest common sub-sequence between x and y , and $\text{len}(x)$ is the length of string x .

Results

We performed experiments using two setups: (1) labeled only, in which the models were trained using only human-labeled data (i.e., the original training data, \mathcal{D}), and (2) synthetic & labeled, in which the models were first trained on synthetic data and then fine-tuned with human-labeled data (i.e., $\tilde{\mathcal{D}} \rightarrow \mathcal{D}$). The validation set was used to determine the best size for the synthetic data based on the F1 score.

Tables 2 and 3 present the experimental results for the general and biomedical domains, respectively. The sequence tagging model generally outperformed the single-span extractive model, which is consistent with the results of previous studies (Segal et al. 2020; Yoon et al. 2022; Li et al. 2022). The CorefRoBERTa model showed robust performance on Quoref because it was designed to capture coreference information. After the best sequence tagging models for each dataset were fine-tuned (i.e., RoBERTa_{large}, CorefRoBERTa_{large}, and BioBERT_{large} for MultiSpanQA, Quoref, and BioASQ, respectively), they outperformed their labeled-only counterparts with F1 scores of 5.0 on MultiSpanQA, 1.9 on Quoref, and 2.8 on the three BioASQ datasets, respectively. In particular, our model outperformed the previous best model on MultiSpanQA (i.e., the multi-task model) by an F1 score of 12.1. For the sequence tagging and single-span extractive models using base-sized encoders, the respective F1 scores improved by 4.5 and 7.1 in the general domain and 3.4 and 9.1 in the biomedical domain, indicating that our framework can be widely utilized with different types of QA models. While the exact-match F1 scores of large models consistently increased across all datasets, the partial-match F1 scores sometimes decreased because of low recall. This

Model	MultiSpanQA	BioASQ 9b
Labeled only	65.0	66.7
<i>Answer extraction methods</i>		
Full Passage	70.3	67.7
Single Sentence (Passage)	70.4	68.1
Single Sentence (Summary)	70.7	68.4
Full Summary [‡]	73.0	71.5
<i>Number of filtering iterations & answer expansion</i>		
w/o Answer expansion		
$T = 0$ (w/o filtering)	71.2	69.1
$T = 1$ (filtering once)	71.3	69.9
$T = 3$ (iterative filtering)	71.6	70.2
+ Answer expansion [‡]	73.0	71.5

Table 4: Ablation study for answer extraction, iterative filtering, and answer expansion. T : maximum number of iterations. [‡]: performance of our final model (LIQUID).

could be because the distribution of the number of answers in our synthetic data is relatively skewed to small numbers compared to that in human-labeled data; thus, models with high capacity might have been biased (Section 4).


4 Analysis

We performed ablation studies of the components and hyperparameters of LIQUID, as well as data quality comparisons of synthetic and human-labeled data. We used 140k synthetic question-answer pairs derived from Wikipedia and PubMed.⁵ For human-labeled data, validation sets of MultiSpanQA and BioASQ 9b were used. Sequence tagging models with RoBERTa_{base} and BioBERT_{base} encoders were used for the list QA model.


Ablation Study

Answer extraction methods We hypothesized that our summarization-based answer extraction method enabled

⁵We will increase the data size up to 1M for each domain and upload them to our official repository.



(a) MultiSpanQA



(b) BioASQ 9b

Figure 3: QA performance for different data sizes of the synthetic data. “Refined data” and “initial data” represent data generated with and without data refinement (i.e., iterative filtering and answer expansion), respectively. The scores at $x = 0$ are obtained using only human-labeled data.

the selection of semantically correlated candidate answers, which are effective for improving list QA performance. To validate this hypothesis, we performed experiments with the following three variants for answer extraction while fixing the iterative filtering and answer expansion methods: (1) full passage, in which all the named entities in the original passage were used as candidate answers; (2) single sentence (passage), in which named entities from a single sentence in the passage were used as candidate answers; and (3) single sentence (summary), in which the passage was summarized, and the named entities were extracted from a single sentence in the summary. Note that our final model was “full summary,” in which a passage was summarized, and the entities from all the sentences in the summary were used. Table 4 indicates that although all three baseline methods improved the performance of the labeled-only models, the performance improvement varied significantly depending on the method. This demonstrates the importance of selecting an appropriate answer extraction method. The full-passage model demonstrated the worst performance, because several unrelated named entities were extracted as candidate answers, as mentioned in Section 1. The single-sentence (passage) and single-sentence (summary) models exhibited similar performance, and both performed better than the full-passage model because the entities in the same single sentence were more likely to be correlated. However, using only the candidate answers within a single sentence could misguide the QG model to create simple questions that do not require complex reasoning over multiple sentences, thereby limiting QA performance (see Section 4). Because our method (i.e., full summary) effectively extracted corre-

Dataset	2	3	4-5	6-9	≥ 10 (%)
Synthetic	76.8	18.3	4.7	0.2	0.0
Labeled	56.8	22.7	13.9	6.1	0.5

Table 5: Number of answer spans in the synthetic and MultiSpanQA (i.e., “labeled”) data.

Answer Type	Synthetic	Labeled
Person	39.6%	39.1%
GPE/LOC	29.8%	18.6%
ORG	19.2%	2.4%
Numeric	1.6%	3.2%
Others	9.7%	19.4%
Non-entity	0.0%	17.4%

Table 6: Distribution of answer types for the synthetic and MultiSpanQA data. GPE/LOC: (non-)geopolitical regions or locations including countries, cities, mountain ranges, etc. ORG: organizations including companies, institutions, sports teams, etc. Others: all other entities in the world. Non-entity: Any phrase that is not defined as an entity.

lated candidate answers from multiple sentences, it significantly outperformed the baselines by F1 scores of 2.3 to 2.7 for the general domain, and 3.1 to 3.8 for the biomedical domain.

Number of filtering iterations & answer expansion We analyzed the effect of filtering by changing the maximum number of filtering steps T . As presented in Table 4, in the absence of filtering (i.e., $T = 0$), we obtained F1 scores of 71.2 and 69.1 for the two domains owing to the noise resulting from incorrect answers. When we used non-iterative filtering (i.e., $T = 1$), we obtained better performance with respective F1 scores of 71.3 and 69.9 for the two domains. Finally, we achieved the best performance with iterative filtering (i.e., $T = 3$) and noted F1 scores of 71.6 and 70.2 for the two domains, respectively. Using more than three steps was not effective because the filtering process was usually completed before $T = 3$. In addition, when we added the answer expansion method, the respective F1 scores for the two domains further improved by 1.4 and 1.3.

Data size We analyzed the variation in the QA performance with the size of the synthetic data for each domain. Figure 3 shows that the performance tended to initially increase and then decrease as the data size increased, indicating the existence of an optimal data size. In addition, the models performed better when iterative filtering and answer expansion methods were used regardless of data size.

Answer Distribution

We analyzed the number of answer spans and answer types in the synthetic data and determined their differences compared to the human-labeled data. Appendix C presents a corresponding analysis of the biomedical domain.

Number of answer spans Table 5 shows that the number of answer spans in the synthetic data tended to be lower than

Question Type	Passage & Answer Spans	Question	Percentage	
			Synthetic	Labeled
Simple Questions	Ya Rab is a 2014 Bollywood movie directed by Hasnain Hyderabadwala starring Ajaz Khan, Arjumman Mughal, Raju Kher, Vikram Singh (actor), Imran Hasnee . . .	Who starred in Ya Rab?	39.3%	26.7%
Lexical Variation	. . . In June 2007, a Hackday event was hosted at Alexandra Palace by the BBC and Yahoo . . .	What <i>media companies</i> hosted a Hackday event in 2007?	60.7%	73.3%
Inter-sentence Reasoning	. . . SBOBET was the shirt sponsor of West Ham United . up until the end of 2012-2013 season. <u>They</u> were also the shirt sponsor of Cardiff City for 2010-2011 season . . .	What teams did SBO-BET sponsor?	33.7%	57.8%
Number of Answers	. . . While working with her mother, Bundy’s uncle offered to pay for her to attend any cookery school in the world. She was accepted into and attended Le Cordon Bleu and Le Notre in Paris, training at Fauchon Patisserie . . .	What <i>two</i> French cookery schools did Bundy attend?	9.0%	7.8%
Entailment	. . . Around the same time, Zhao Yun also came to Ye (present-day Handan, Hebei), Yuan Shao’s headquarters, where he met Liu Bei again . . .	Who were the people who came to Ye?	1.1%	3.3%

Table 7: Classification of questions in the synthetic and MultiSpanQA data. All examples are from the synthetic data. Answers are represented in bold. See the main text for descriptions of the question types.

that in MultiSpanQA. The majority (76.8%) of the questions had two spans, but certain questions (4.9%) had more than three answers. We aim to perform further analyses to determine whether dataset bias is caused by limited answer spans.

Answer types We manually classified the types of answers to 100 questions into entity-type categories (Table 6). Notably, both synthetic and labeled datasets consist of many person and geopolitical/location entities. Apart from these two answer types, humans seem to have a tendency to ask various questions that are not limited to particular answer types. This results in relatively few organization-type answers and several others-type answers. The most notable difference between the two datasets is the number of non-entity answers. As we relied on NER to extract answers, we could not effectively deal with answers that differed from the named entities defined in the NER system, which is a limitation of our framework.⁶

Question Types

We explored the quality and types of list questions generated by the model and asked by humans. We sampled 100 questions from the synthetic and MultiSpanQA data but used only 89 and 90 correct examples for the model- and human-generated questions, respectively, after excluding the incorrect examples (see Appendix C for a corresponding analysis in the biomedical domain). We manually classified the questions into the following five categories based on the reasoning required to answer these questions:

- **Simple questions:** Questions that were simply derived from evidence texts with few lexical variations.

⁶The QA model can extract non-entity answers, but we did not find such questions in the 100 sampled examples.

- **Lexical variation:** Questions that were created with lexical variations using synonyms and hypernyms.
- **Inter-sentence reasoning:** Questions that required high-level reasoning such as anaphora, or answers that were distributed across multiple sentences.
- **Number of answers:** Questions that specified the number of answers, which is a characteristic of list questions.
- **Entailment:** Questions that required textual entailment based on the evidence texts and commonsense.

Table 7 lists data examples generated by LIQUID and the proportion of each question type in the synthetic and labeled data. It is worth noting that, although the QG model was not tuned for list-type questions, the resulting questions require various types of reasoning to obtain multiple correct answers. Numerous questions contained lexical variations (60.7%) while questions that simply matched the evidence texts were avoided. Some questions involved inter-sentence reasoning (33.7%) or textual entailment (1.1%) or specified the number of answer spans (9.0%).

Additionally, we discovered that most simple questions were generated when the answers belonged to single sentences, indicating the importance of extracting answers spread across multiple sentences to create complex questions. This also indicates that because the single-sentence models (Table 4) extracted candidate answers from single sentences, their performance was relatively low compared to that of LIQUID, which extracted answers from multiple relevant sentences within the summary.

Error Analysis

We analyzed 411 synthetic data examples and discovered that 50 of them (12.2%) were wrong.⁷ The most dominant error type involved the presence of incorrect answer spans, accounting for 78% of all errors. These errors occurred when unrelated entities were grouped into the same set of candidate answers during the answer extraction process, and the filtering model failed to eliminate them. This indicates that achieving a high accuracy in the answer set still remains a challenge. In addition, 12% of the errors can be attributed to missing answers not detected by the expansion model, and 4% of the errors are incorrect answers added by the model, indicating that the answer expansion method must be improved in future studies. Finally, the remaining errors (8%) can be attributed to the QG model, which can be reduced by developing large-scale list QA data or pre-trained models.

5 Conclusion

Herein, we introduced LIQUID, a framework that automatically generates list QA datasets from unlabeled corpora to alleviate the data scarcity problem in this field. Our synthetic data significantly improved the performance of the current supervised models on five benchmark datasets. We thoroughly analyzed the effect of each component in LIQUID and generated data quantitatively and qualitatively.

A Algorithm

Algorithm 1 describes the data generation process in detail.

B Implementation Details

LIQUID

For the unlabeled corpus, we used the 2018-12-20 version of Wikipedia and the 2019-01-02 version of PubMed abstracts for the general and biomedical domains, respectively. We re-used the trained model parameters available online for the summarization,⁸ QG,⁹ and QA models.^{10,11} The minimum and maximum lengths of the output sequence were set to 64 and 128 for the summarization model and 32 and 128 for the QG model. For the QA model, the maximum lengths of the question and evidence text were set to 128 and 384, respectively, and the trained checkpoints of the RoBERTa_{base}¹² and BioBERT_{base}¹³ models were used. We manually selected the thresholds τ , 0.1, and 0.05 for the general and biomedical domains, respectively. In NER, we excluded the

⁷In MultiSpanQA, 10% of the examples appear to be incorrect, indicating that our framework can generate QA data as accurately as human annotators. Nevertheless, humans are superior in terms of data diversity and quality, as described in Sections 4 and 4.

⁸huggingface.co/facebook/bart-large-cnn

⁹huggingface.co/mrm8488/t5-base-finetuned-question-generation-ap

¹⁰huggingface.co/thatdramebaazguy/roberta-base-squad

¹¹huggingface.co/dmis-lab/biobert-base-cased-v1.1-squad

¹²huggingface.co/thatdramebaazguy/roberta-base-squad

¹³huggingface.co/dmis-lab/biobert-base-cased-v1.1-squad

Algorithm 1: Data generation process of LIQUID

Input: Source corpus \mathcal{C} ; Summarization model θ_S ; NER model θ_{NER} ; QG model θ_{QG} ; (single-span) QA model θ_{QA} ; **Parameter:** The number of passages to sample K ; The number of entity types L ; The number of filtering iterations T ; Threshold for iterative filtering and answer expansion τ ; **Output:** Synthetic QA dataset $\tilde{\mathcal{D}}$;

```
1:  $\tilde{\mathcal{D}} \leftarrow \{\}$ 
2: for  $k \leftarrow 1$  to  $K$  do
3:    $c_k \leftarrow \text{RandomSampling}(\mathcal{C})$ 
4:    $\tilde{c}_k \leftarrow \text{Summarization}(c_k; \theta_S)$ 
5:    $\mathbf{A}_1, \dots, \mathbf{A}_L \leftarrow \text{AnswerExtraction}(\tilde{c}_k; \theta_{\text{NER}})$ 
6:   /* Assume that every answer set with less than two
   elements is excluded from the process. */
7:   for  $l \leftarrow 1$  to  $L$  do
8:      $q_l \leftarrow \text{QuestionGeneration}(\mathbf{A}_l, c_k; \theta_{\text{QG}})$ 
9:      $\mathbf{A}_l^0, \mathbf{A}_l^1, q_l^1 \leftarrow \{\}, \mathbf{A}_l, q_l$ 
10:     $t \leftarrow 1$ 
11:    while  $\mathbf{A}_l^{t-1} \neq \mathbf{A}_l^t$  &  $t \leq T$  do
12:       $\mathbf{A}_l^{t+1} \leftarrow \text{Filtering}(c_k, q_l^t, \mathbf{A}_l^t; \theta_{\text{QA}}, \tau)$ 
13:       $q_l^{t+1} \leftarrow \text{QuestionGeneration}(\mathbf{A}_l^{t+1}, c_k; \theta_{\text{QG}})$ 
14:       $t \leftarrow t + 1$ 
15:    end while
16:     $\mathbf{A}_l', q_l' \leftarrow \mathbf{A}_l^{t-1}, q_l^{t-1}$ 
17:     $\mathbf{A}_l'' \leftarrow \text{AnswerExpansion}(c_k, q_l', \mathbf{A}_l'; \theta_{\text{QA}}, \tau)$ 
18:     $q_l'' \leftarrow \text{QuestionGeneration}(\mathbf{A}_l'', c_k; \theta_{\text{QG}})$ 
19:     $\tilde{\mathbf{A}}_l \leftarrow \text{Filtering}(c_k, q_l'', \mathbf{A}_l''; \theta_{\text{QA}}, \tau)$ 
20:    if  $\mathbf{A}_l'' = \tilde{\mathbf{A}}_l$  then
21:       $\tilde{\mathcal{D}} \leftarrow \text{Append}(\tilde{\mathcal{D}}, (c_k, q_l'', \tilde{\mathbf{A}}_l))$ 
22:    else
23:       $\tilde{\mathcal{D}} \leftarrow \text{Append}(\tilde{\mathcal{D}}, (c_k, q_l', \mathbf{A}_l''))$ 
24:    end if
25:  end for
26: end for
```

date and species types for the general and biomedical domains, respectively, because they led to trivial candidate answers in our initial experiments. The best data sizes and best validation F1 scores are listed in Table 8.

List QA models

We implemented single-span extractive models by modifying the BioBERT-PyTorch repository.¹⁴ Sequence tagging models were implemented using the code provided by Yoon et al. (2022). For both types of models, we used a batch size of 8 and learning rate of 1e-4. We searched for the best threshold values of the single-span extractive models using the F1 score on the validation set. As a result of searching from 0.02 to 0.15 in 0.01 intervals, we selected 0.1 (Quoref), 0.03 (BioASQ 7b), 0.04 (BioASQ 8b), and 0.1 (BioASQ 9b). For MultiSpanQA, we used a dynamic threshold search method following Li et al. (2022).

¹⁴github.com/dmis-lab/biobert-pytorch

Model	Data Size $ \tilde{\mathcal{D}} $	Validation F1
<i>MultiSpanQA, Quoref</i>		
RoBERTa _{base} + S	60k, 60k	22.3, 60.7
RoBERTa _{base} + T	30k, 90k	73.0, 85.7
RoBERTa _{large} + T	50k, 60k	73.6, 86.7
CorefRoBERTa _{large} + T	40k, 20k	70.3, 88.4
<i>BioASQ 7b, 8b, 9b</i>		
BioBERT _{base} + S	80k, 130k, 80k	51.8, 44.2, 64.0
BioBERT _{base} + T	80k, 80k, 60k	49.0, 44.2, 71.5
BioBERT _{large} + T	40k, 40k, 60k	51.5, 46.3, 72.2

Table 8: Best synthetic data sizes and corresponding F1 scores on the validation sets for each benchmark dataset.

Dataset	2	3	4-5	≥ 6 (%)
Synthetic	61.6	24.2	13.2	1.0
BioASQ 9b	44.3	26.5	26.5	2.8

Table 9: Comparison of the number of answer spans.

Answer Type	Synthetic	BioASQ 9b
Disease	47.6%	29.6%
Drug/Chemical	29.5%	21.8%
Gene/Protein	14.2%	28.9%
Cell Type/Cell Line	4.7%	0.0%
Others	0.8%	14.8%
Non-entity	3.1%	4.9%

Table 10: Comparison of answer types.

Datasets and Evaluation

We obtained MultiSpanQA, Quoref, and BioASQ 9b from the official websites. For BioASQ 7b and 8b, we used document-level evidence texts provided from the repository of Yoon et al. (2022). We used the evaluation script in the official repository of Li et al. (2022).

C Biomedical Data Analysis

We analyzed synthetic and human-labeled data for the biomedical domain using 140k question-answer pairs generated from PubMed and the BioASQ 9b validation set.

Number of answer spans Table 9 presents the distribution of the number of answers. Similar to the results in the general domain (Table 5), the synthetic data were more skewed toward smaller numbers of answers than the labeled data, but some answers (14.2%) had four or more spans.

Answer types We analyzed the answer types of 100 and 50 examples in the synthetic and labeled data, respectively. As shown in Table 10, the disease, drug/chemical, and gene/protein types were dominant in both datasets. BioASQ 9b does not seem to contain cell types and cell lines and consists of many others-type answers, mainly organs. The synthetic data contains some non-entity answers,

Question Type	Synthetic	BioASQ 9b
Simple Question	58.2%	22.2%
Lexical Variation	40.7%	51.1%
Domain Knowledge	1.1%	26.7%
Inter-sentence Reasoning	8.8%	13.3%
Number of Answers	29.7%	15.6%

Table 11: Comparison of question types.

Stage	Required Time	
	Wikipedia	PubMed
1. Answer Extraction		
- Summarization	26m 56s	26m 19s
- NER	1m 43s	36m 2s
2. Question Generation	9m 21s	4m 39s
3. Iterative Filtering & Answer Expansion	33m 20s	20m 23s
Total Time	1h 11m 20s	1h 27m 23s

Table 12: The time taken to process 10k passages.

such as descriptions of symptoms and treatments, which are added in the answer expansion stage.

Question types We classified 91 and 45 correct examples of 100 and 50 examples for the synthetic and labeled data, respectively, excluding incorrect examples. Unlike for the general domain (Table 7), we did not use the entailment category and added the domain knowledge category, where questions required some biomedical knowledge. As shown in Table 11, many of the synthetic questions contained lexical variations (40.7%). The QG model sometimes generates questions that require domain knowledge, but at a much lower rate than the data annotated by human experts.

D Efficiency

We measured the time required to process 10k passages from Wikipedia and PubMed. We ran our model on an Intel(R) Xeon(R) Silver 4210R CPU @ 2.40GHz and a single 24GB GPU (GeForce RTX 3090). We used a batch size of eight, that is, eight passages were processed simultaneously. Table 12 shows that we can process 10k Wikipedia passages in 72 minutes and 10k PubMed passages in 88 minutes. For each corpus, we obtained 8,950 and 5,190 initial questions, and 4,274 (47.8%) and 2,654 (51.1%) questions after the iterative filtering and answer expansion stages, respectively. In the question generation, iterative filtering, and answer expansion stages, the processing of passages in Wikipedia took more time than in PubMed because the number of entity types used in the general domain (17 types) was approximately twice than that in the biomedical domain (8 types). However, the overall process with PubMed was relatively slow mainly because of the run time of the NER model.

Acknowledgements

We thank Wonjin Yoon, Gangwo Kim, and Hajung Kim for the helpful feedback on our manuscript. We thank Haonan Li for helping us evaluate our models on the MultiSpanQA leaderboard. This research was supported by (1) National Research Foundation of Korea (NRF-2020R1A2C3010638), (2) the MSIT (Ministry of Science and ICT), Korea, under the ICT Creative Consilience program (IITP-2021-2020-0-01819) supervised by the IITP (Institute for Information & communications Technology Planning & Evaluation), and (3) a grant of the Korea Health Technology R&D Project through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea (grant number: HR20C0021).

References

- Alberti, C.; Andor, D.; Pitler, E.; Devlin, J.; and Collins, M. 2019. Synthetic QA Corpora Generation with Roundtrip Consistency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6168–6173. Florence, Italy: Association for Computational Linguistics.
- Dasigi, P.; Liu, N. F.; Marasović, A.; Smith, N. A.; and Gardner, M. 2019. Quoref: A Reading Comprehension Dataset with Questions Requiring Coreferential Reasoning. 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)*, 5925–5932. Hong Kong, China: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*.
- Dhingra, B.; Danish, D.; and Rajagopal, D. 2018. Simple and Effective Semi-Supervised Question Answering. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 582–587. New Orleans, Louisiana: Association for Computational Linguistics.
- Joshi, M.; Choi, E.; Weld, D.; and Zettlemoyer, L. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1601–1611. Vancouver, Canada: Association for Computational Linguistics.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; Toutanova, K.; Jones, L.; Kelcey, M.; Chang, M.-W.; Dai, A. M.; Uszkoreit, J.; Le, Q.; and Petrov, S. 2019. Natural Questions: A Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics*, 7: 452–466.
- Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C. H.; and Kang, J. 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4): 1234–1240.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. 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*, 7871–7880. Online: Association for Computational Linguistics.
- Lewis, P.; Wu, Y.; Liu, L.; Minervini, P.; Küttler, H.; Piktus, A.; Stenetorp, P.; and Riedel, S. 2021. PAQ: 65 Million Probably-Asked Questions and What You Can Do With Them. *Transactions of the Association for Computational Linguistics*, 9: 1098–1115.
- Li, H.; Tomko, M.; Vasardani, M.; and Baldwin, T. 2022. MultiSpanQA: A Dataset for Multi-Span Question Answering. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1250–1260. Seattle, United States: Association for Computational Linguistics.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lu, K.; Hsu, I.; Zhou, W.; Ma, M. D.; Chen, M.; et al. 2022. Summarization as Indirect Supervision for Relation Extraction. *arXiv preprint arXiv:2205.09837*.
- Lyu, C.; Shang, L.; Graham, Y.; Foster, J.; Jiang, X.; and Liu, Q. 2021. Improving Unsupervised Question Answering via Summarization-Informed Question Generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 4134–4148. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Nallapati, R.; Zhou, B.; dos Santos, C.; Gulcehre, C.; and Xiang, B. 2016. Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, 280–290. Berlin, Germany: Association for Computational Linguistics.
- Nentidis, A.; Bougiatiotis, K.; Krithara, A.; and Paliouras, G. 2019. Results of the seventh edition of the BioASQ challenge. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 553–568. Springer.
- Nentidis, A.; Katsimpras, G.; Vandorou, E.; Krithara, A.; Gasco, L.; Krallinger, M.; and Paliouras, G. 2021. Overview of bioasq 2021: The ninth bioasq challenge on large-scale biomedical semantic indexing and question answering. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, 239–263. Springer.
- Nentidis, A.; Krithara, A.; Bougiatiotis, K.; Krallinger, M.; Rodriguez-Penagos, C.; Villegas, M.; and Paliouras, G. 2020. Overview of bioasq 2020: The eighth bioasq challenge on large-scale biomedical semantic indexing and ques-

- tion answering. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, 194–214. Springer.
- Puri, R.; Spring, R.; Shoeybi, M.; Patwary, M.; and Catanzaro, B. 2020. Training Question Answering Models From Synthetic Data. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 5811–5826. Online: Association for Computational Linguistics.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P. J.; et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140): 1–67.
- Rajpurkar, P.; Jia, R.; and Liang, P. 2018. Know What You Don't Know: Unanswerable Questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 784–789. Melbourne, Australia: Association for Computational Linguistics.
- Rajpurkar, P.; Zhang, J.; Lopyrev, K.; and Liang, P. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2383–2392. Austin, Texas: Association for Computational Linguistics.
- Segal, E.; Efrat, A.; Shoham, M.; Globerson, A.; and Berant, J. 2020. A Simple and Effective Model for Answering Multi-span Questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3074–3080. Online: Association for Computational Linguistics.
- Shakeri, S.; Nogueira dos Santos, C.; Zhu, H.; Ng, P.; Nan, F.; Wang, Z.; Nallapati, R.; and Xiang, B. 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)*, 5445–5460. Online: Association for Computational Linguistics.
- Sung, M.; Jeong, M.; Choi, Y.; Kim, D.; Lee, J.; and Kang, J. 2022. BERN2: an advanced neural biomedical named entity recognition and normalization tool. *Bioinformatics*, 38(20): 4837–4839.
- Trischler, A.; Wang, T.; Yuan, X.; Harris, J.; Sordani, A.; Bachman, P.; and Suleman, K. 2017. NewsQA: A Machine Comprehension Dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, 191–200. Vancouver, Canada: Association for Computational Linguistics.
- Tsatsaronis, G.; Balikas, G.; Malakasiotis, P.; Partalas, I.; Zschunke, M.; Alvers, M. R.; Weissenborn, D.; Krithara, A.; Petridis, S.; Polychronopoulos, D.; et al. 2015. An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC bioinformatics*, 16(1): 1–28.
- Voorhees, E. M.; et al. 2001. Overview of the TREC 2001 Question Answering Track. In *Trec*, 42–51.
- Yang, Z.; Hu, J.; Salakhutdinov, R.; and Cohen, W. 2017. Semi-Supervised QA with Generative Domain-Adaptive Nets. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1040–1050. Vancouver, Canada: Association for Computational Linguistics.
- Yang, Z.; Qi, P.; Zhang, S.; Bengio, Y.; Cohen, W.; Salakhutdinov, R.; and Manning, C. D. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2369–2380. Brussels, Belgium: Association for Computational Linguistics.
- Ye, D.; Lin, Y.; Du, J.; Liu, Z.; Li, P.; Sun, M.; and Liu, Z. 2020. Coreferential Reasoning Learning for Language Representation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 7170–7186. Online: Association for Computational Linguistics.
- Yoon, W.; Jackson, R.; Lagerberg, A.; and Kang, J. 2022. Sequence tagging for biomedical extractive question answering. *Bioinformatics*, 38(15): 3794–3801.