

Multi-modal Retrieval of Tables and Texts Using Tri-encoder Models

Bogdan Kostić and Julian Risch and Timo Möller

deepset

{bogdan.kostic, julian.risch, timo.moeller}@deepset.ai

Abstract

Open-domain extractive question answering works well on textual data by first retrieving candidate texts and then extracting the answer from those candidates. However, some questions cannot be answered by text alone but require information stored in tables. In this paper, we present an approach for retrieving both texts and tables relevant to a question by jointly encoding texts, tables and questions into a single vector space. To this end, we create a new multi-modal dataset based on text and table datasets from related work and compare the retrieval performance of different encoding schemata. We find that dense vector embeddings of transformer models outperform sparse embeddings on four out of six evaluation datasets. Comparing different dense embedding models, tri-encoders with one encoder for each question, text and table increase retrieval performance compared to bi-encoders with one encoder for the question and one for both text and tables. We release the newly created multi-modal dataset to the community so that it can be used for training and evaluation.

1 Introduction

Finding the answer to a factual question in a large collection of documents is a tedious task that many people from a broad range of domains have to complete on a daily basis. In order to address this task with machine learning approaches, it has been formalized as open-domain extractive question-answering (QA). More specifically, given a natural language question and a database of text documents as a knowledge base, open-domain extractive QA aims to extract a substring that answers the given question out of one of the documents.

The standard approach for this task is a pipeline architecture consisting of two components: a retriever selecting a small subset of relevant documents from the database and a reader extracting granular answers out of each of these retrieved documents (Voorhees and Tice, 2000). In this paper,

we focus on the retriever and present a transformer-based tri-encoder model as an implementation of this component. Retrievers can also be implemented with bag-of-words retrieval methods, such as TF-IDF or BM25, that transform all documents and the question to sparse vector representations. However, as these methods rely on a lexical overlap of the question and the documents, they fail to capture synonymy and other semantic relationships. This limitation motivates the use of dense vector representations and we compare dense retrieval models to a BM25 baseline in our experiments. A survey on term-based, early semantic, and neural semantic models for document retrieval as a first step before document reranking and down-stream tasks, such as question answering, has been published by Cai et al. (2021).

To date, most of the research centered around question-answering focuses on using free-form text as the single source for answering questions. However, valuable information can obviously be found in other modalities as well. For instance, a lot of information is stored in semi-structured tables; according to Cafarella et al. (2008), more than 14.1 billion tables can be found on the World Wide Web. Given that a user typically does not know in advance in which modality the answer to their question resides, a QA system capable of jointly handling text and tables is needed. One major challenge in building such a system is to represent texts and tables in a way that allows capturing semantic similarity and retrieving texts and tables that are semantically related to a given question.

Contributions The contributions of this paper can be summarized as follows: (1) we present bi-encoder and tri-encoder models that are capable of joint retrieval of tables and texts; (2) we create and release a multi-modal dataset for training and

evaluating models on this task;¹ (3) we compare sparse retrieval models with dense retrieval models using bi-encoders and tri-encoders on our new multi-modal dataset and on five uni-modal datasets from related work.

Outline The remainder of this paper is structured as follows: Section 2 summarizes existing methods for uni-modal retrieval of texts on the one hand and of tables on the other hand. Further, it discusses the only two, recently published approaches for joint retrieval of tables and texts. Section 3 briefly describes the existing uni-modal datasets and our new multi-modal dataset, which we use to train the retrieval models presented in Section 4 and to evaluate these models in Section 5. Section 6 concludes the paper and gives directions for future work.

2 Related Work

To the best of our knowledge the only existing work that addresses joint retrieval of tables and texts is by Chen et al. (2021), Talmor et al. (2021) and Li et al. (2021). To address the challenge of limited context available for tables, Chen et al. (2021) fuse a table segment and text passages into one block if they mention the same named entities. Each block is represented with a single dense embedding so that a table and relevant passages are jointly retrieved as a group if the embedding is similar to that of the question. This grouping makes sense because Chen et al. (2021) address the task of multi-hop QA, where information needs to be aggregated from multiple tables and texts to answer a question. In contrast to that, we address the slightly different task of single-hop QA, where either only one table or one text is needed to answer a question. Therefore, our approach represents tables and texts with separate embeddings, which are in the same embedding space. The advantage here is that the model learns to estimate relevance on the more fine-grained level of individual tables or texts and can decide whether a particular table or text is more relevant to the question.

Li et al. (2021) also address the task of multi-hop QA on texts and tables but retrieve them individually. They make use of the sparse retrieval method BM25 and a transformer-based reranker to generate a set of candidate texts and tables. With two separate BM25 indices for texts and tables, they retrieve

a set of documents for each modality. In a second step, they apply a joint BERT-based reranker to reduce the size of candidate texts and tables. Talmor et al. (2021) create a MULTIMODALQA dataset containing questions that require joint reasoning over tables, texts, and images.

Otherwise closely related to our approach are TABERT by Yin et al. (2020) and TAPAS by Herzig et al. (2020), who use tables and texts for language model pre-training but do not consider joint retrieval of tables and texts. TURL focuses on representation learning for tables but also does not consider the retrieval task (Deng et al., 2020). Due to this limited amount of prior research, we discuss related work on the separate tasks of uni-modal text retrieval and table retrieval in the following.

2.1 Text Retrieval

Dense Passage Retrieval (DPR) by Karpukhin et al. (2020) relies on a bi-encoder model comprising two separate BERT models (Devlin et al., 2019). Similar to TF-IDF, DPR is a vector space model that represents both queries and documents in the same vector space. However, while TF-IDF represents text documents as very high dimensional sparse vectors, DPR relies on relatively low dimensional dense embeddings. While one of the models, the passage encoder ($BERT_p$), is used to encode text passages at indexing time, the second model, the question encoder ($BERT_q$), is used to encode questions at query time. Since BERT’s [CLS]-token is particularly designated to capture the meaning of the whole input sequence, its embedding is used as a representation vector for both the text passages and the questions.

The training aims to increase the dot product or cosine similarity of semantically similar passages and questions. In order to achieve this, Karpukhin et al. (2020) use, besides the question’s positive passage, hard negative passages as well as in-batch negative passages as training signal. Hard negatives are sampled utilizing a BM25-based retriever on the whole English Wikipedia dump. For each question, they use the highest ranked passage not containing the question’s answer string. DPR drastically outperforms BM25 by almost 20 percentage points with regard to recall@20 on the Natural Questions (NQ) dataset by Kwiatkowski et al. (2019).

¹<https://multimodalretrieval.s3.eu-central-1.amazonaws.com/data.zip>

2.2 Table Retrieval

Most existing table retrieval approaches rely on supervised learning-to-rank approaches. [Zhang and Balog \(2018\)](#) combine a set of hand-crafted query features, table features and query-table features with semantic similarity of table and query as additional feature. To get table and query representations, they use the average of pre-trained word2vec ([Mikolov et al., 2013](#)) and RDF2vec embeddings ([Ristoski and Paulheim, 2016](#)). These features are then used to train a random forest regressor to get relevance scores of the tables with regard to the query. Interestingly, training word embeddings on a corpus of tables instead of texts does not improve performance ([Zhang et al., 2019](#)).

[Shraga et al. \(2020a\)](#) combine intrinsic and extrinsic table similarity scores. For the intrinsic table similarity, they concatenate the table’s title, caption, header and data rows and use a sliding window over this concatenated string to get a set of candidate passages for each table. Each candidate passage is scored using BM25 and the maximum of all passages of a table is the intrinsic score with regard to a query. The extrinsic score is based on table-table similarities to ensure that similar tables get a similar final score.

[Bagheri and Al-Obeidat \(2020\)](#) focus on hard queries that contain terms that do not occur in the relevant tables, which means the query and the relevant tables have a low lexical overlap. To this end, they learn low dimensional latent factor matrices to represent tables as well as queries, i.e., they learn term co-occurrences to be able to get tables that address the same topic but only partially overlap on a lexical level. Based on this result, we compare results on datasets with high or low lexical overlap in our experiments.

There are four deep learning approaches for table retrieval ([Shraga et al., 2020b](#); [Pan et al., 2021](#); [Chen et al., 2020b](#); [Herzig et al., 2021](#)). [Shraga et al. \(2020b\)](#) treat tables and queries as multi-modal objects that consist of query, table caption, schema, rows and columns. Each component is encoded using its own neural network encoder that accounts for its special characteristics. Subsequently, these uni-modal encodings are joined into a single representation, which is passed on to fully connected layers that predict whether the input table is relevant with regard to the input query.

[Pan et al. \(2021\)](#) stack two retrieval components. First, they use BM25 to produce a large subset

of possibly relevant tables. These tables are then passed to a row-column intersection model that generates a probability distribution over the table cells whether they contain the answer to the user’s query. The maximum cell-level score for each table represents its retrieval score.

[Chen et al. \(2020b\)](#) apply the transformer-based language model BERT ([Devlin et al., 2019](#)) to the table retrieval task by combining BERT embeddings with other, hand-curated table and query features. Their approach consists of several components, including a step where the concatenation of a query, a table’s context fields and selected relevant table rows are processed by a BERT model. A significant downside of this approach is that it becomes inefficient with increasing number of tables: for each query, all tables need to be passed through a BERT network.

[Herzig et al. \(2021\)](#) solve this efficiency problem by adapting [Karpukhin et al.’s \(2020\)](#) dense passage retrieval approach to dense table retrieval (DTR). To this end, they make use of TAPAS ([Herzig et al., 2020](#)), a transformer-based language model that has been pre-trained on millions of tables. TAPAS extends BERT by adding three different types of positional embeddings to encode the two-dimensional tabular structure: row, column and rank embeddings. This allows to flatten the table by concatenating the rows to a one-dimensional sequence of tokens. Similar to [Karpukhin et al. \(2020\)](#), [Herzig et al. \(2021\)](#) make use of a bi-encoder approach. However, they use two TAPAS instances instead of BERT instances to encode the queries and the tables, respectively. The goal of training this bi-encoder is to build an embedding model that generates similar embeddings for questions and their relevant tables. As in [Karpukhin et al.’s \(2020\)](#) approach, this goal is achieved using hard-negatives retrieved from all the tables from the English Wikipedia dump as well as in-batch negatives. DTR outperforms BM25 by more than 40 percentage points on the NQ-TABLES dataset ([Herzig et al., 2021](#)). However, the experiments also show that TAPAS requires additional pre-training on the task of table retrieval on millions of tables scraped from Wikipedia. As a further research direction for future work, [Herzig et al. \(2021\)](#) propose to combine tables and texts for multi-modal open-domain QA. We contribute towards this goal in our paper by providing a multi-modal retriever as one component of a multi-modal

open-domain QA pipeline on tables and texts.

3 Datasets

The training and evaluation of models examined in this paper on the task of multi-modal retrieval makes use of five datasets from related work: NQ (Kwiatkowski et al., 2019), NQ-TABLES (Herzig et al., 2021), WIKISQL (Zhong et al., 2017), a subset of WIKISQL, which we call WIKISQL_{ctx-independent}, and OTT-QA (Chen et al., 2021). This section briefly explains the characteristics of these datasets and of our newly created multi-modal retrieval dataset comprising tables and texts, which we call MULTIMODALRETRIEVAL. Table 1 gives an overview of the modality and the number of samples in each dataset.

NQ Natural Questions (NQ) (Kwiatkowski et al., 2019) is an open-domain QA dataset on Wikipedia articles. It consists of questions, their answers and the text passages the answers reside in. The questions are *natural* because they consist of real user queries issued to the Google search engine instead of questions posed by annotators after reading a text passage, which was done to create other popular QA datasets, such as the Stanford Question Answering Dataset (SQuAD) (Rajpurkar et al., 2018). Having *natural* questions does not only ensure that the questions correspond to information needs by real users but also makes the dataset open-domain, i.e., the questions are context-independent and can be understood without their accompanying text passage that contains the answer. For the purpose of retrieval, we utilize Karpukhin et al.’s (2020) pre-processed variant of NQ.

NQ-TABLES The answers to most of the questions inside the NQ dataset can be found in plain unstructured text. However, a subset of the questions are answered by tables. As a consequence, Herzig et al. (2021) construct NQ-TABLES, a table-specific QA dataset based on NQ. To achieve this objective, they extract all the tables and all the questions whose answers reside inside a table. They come up with a dataset consisting of 9,594 questions in the training set, 1,068 questions in the development set, 966 questions in the test set and 169,898 tables in total. Given that this dataset is a subset of NQ, the questions share the characteristic of being context-independent.

WIKISQL WIKISQL (Zhong et al., 2017) is a closed-domain QA dataset consisting of 24,241 ta-

bles and 80,654 natural language questions together with their corresponding SQL query. To build this dataset, Zhong et al. (2017) generate a number of random SQL queries for each table. These SQL queries are then transformed into crude questions using templates. Finally, Amazon Mechanical Turk crowd workers paraphrase these crude questions into natural language questions, which are checked by two additional crowd workers.

WIKISQL_{ctx-independent} Since WIKISQL is a closed-domain QA dataset, the majority of its questions is context-dependent, i.e., they do not provide enough context to be answered without their accompanying table and are therefore not suitable for training or evaluation on the retrieval task. One example for such insufficient questions inside the dataset is: “*Who is the player that wears number 42?*”, which cannot be answered without additional context given in the table, such as the name of a sports team and a year. As a consequence, to make use of WIKISQL, all questions that do not provide enough context for retrieval need to be filtered out. For automating this filtering, we labeled a subset of WIKISQL’s questions with regard to whether they are either context-independent or under-specified resulting in 4,553 labels as training set and 612 labels as test set. These labels are then used to train a classifier that predicts whether a question provides enough context. We fine-tune a RoBERTa-base (Liu et al., 2019) language model achieving an accuracy of 0.8134 and a macro-averaged F1-score of 0.7748 on the test set. Next, we apply this classifier to the whole WIKISQL dataset to filter out all the questions that are predicted as under-specified.

OTT-QA OTT-QA (Chen et al., 2021) is an open-domain multi-hop QA dataset of texts and tables from Wikipedia built upon HybridQA (Chen et al., 2020a), a closed-domain multi-hop QA dataset. Multi-hop QA refers to the fact that most questions inside the dataset require a combination of different texts and/or tables instead of a single document in order to be answered. To generate an open-domain version of HybridQA, Chen et al. (2021) let crowd workers decontextualize all the questions. Furthermore, they add additional question-answer pairs on newly crawled tables. Since the published annotations contain only the gold tables but not the gold texts, we use OTT-QA only for generating table retrieval training samples and evaluating the uni-modal retrieval of tables.

Dataset	Modality	Train	Test	Ctx-ind.
NQ	text	58,880	3,610	✓
NQ-TABLES	table	9,594	966	✓
WIKISQL	table	56,355	15,878	✗
WIKISQL _{ctx-independent}	table	7,336	2,101	✓
OTT-QA	table	41,469	2,158	✓
MULTIMODALRETRIEVAL	text & table	120,239	4,937	✓

Table 1: Modality and number of train and test samples for multi-modal retrieval models. Only WIKISQL is not context-independent (ctx-ind.), which is why we create the subset WIKISQL_{ctx-independent}.

MULTIMODALRETRIEVAL Given that no multi-modal dataset of tables and texts is readily available, this paper newly introduces such a dataset based on datasets from related work. For this purpose, we combine the question-passage pairs from NQ as questions requiring a text passage to be answered with the question-table pairs from NQ-TABLES, WIKISQL_{ctx-independent}, and OTT-QA as questions requiring a table to be answered. Both Karpukhin et al. (2020) for text retrieval and Herzig et al. (2021) for table retrieval show that adding hard negatives as training signal boosts the retrieval performance of the model significantly. Therefore, we index 21 million Wikipedia passages (from Karpukhin et al. (2020)) and 7 million Wikipedia tables (used by Eisenschlos et al. (2020) to pre-train TAPAS) with Elasticsearch to sample hard negatives using BM25. For each question, the highest ranked passages or tables that do not contain the answer string were chosen, i.e., a question originating from a tabular question-answering dataset can also have a text passage as hard negative, and vice versa.

4 Approach

Our two approaches for the joint retrieval of tables and texts comprise bi-encoder and tri-encoder models based on uni-modal dense retrieval methods by Karpukhin et al. (2020) and Herzig et al. (2021).

In our first approach, the bi-encoder uses one language model to encode the questions and a second model to encode both the tables and the text passages. Our second approach adds a third encoder, such that there is one separate encoder for questions, text passages and tables. In contrast to the bi-encoder where tables and texts are encoded by the same model, the tri-encoder routes tables to the table encoder model and text passages to the

text encoder model.

We train three different bi-encoders and three different tri-encoders, which differ in the underlying language models as specified in Table 2. The first multi-modal bi-encoder consists of two different BERT-small instances that serve as question encoder and table and text encoder, respectively. Given that BERT models only allow one-dimensional strings of text as input, the two-dimensional tables are transformed into one dimension by concatenating the titles of the page and the section the table occurs in, the caption of the table, and each row of the table. In order to analyze whether the table-specific language model TAPAS (Herzig et al., 2020; Eisenschlos et al., 2020) gives a performance boost compared to a plain BERT model, the remaining two bi-encoders make use of a TAPAS model that is pre-trained for the task of table retrieval (Herzig et al., 2021) for at least one of their encoders. Thus, the second bi-encoder uses a BERT-small instance as question encoder and a TAPAS-small instance as table and text encoder. The third bi-encoder utilizes two TAPAS-small models, one to encode the questions and the second to encode both text and tables. Using BERT-small instead of BERT-base or BERT-large models drastically reduces the number of parameters and allows to fit more training samples into one batch. In contrast to a BERT-large model with 24 transformer layers, hidden representations of size 1024, 16 attention heads, and a total number of 335M parameters, BERT-small consists of only 4 transformer layers, hidden representations of size 512, 8 attention heads, and a total number of 29.1M parameters.

The first tri-encoder model uses BERT-small instances for all of its three encoders. Also for the tri-encoder approach, we analyze the impact of using TAPAS for at least one of the encoders. Therefore,

		Encoders	
	Question	Text	Table
Bi-encoder	BERT-small	—— BERT-small ——	
	BERT-small	—— TAPAS-small ——	
	TAPAS-small	—— TAPAS-small ——	
Tri-encoder	BERT-small	BERT-small	BERT-small
	BERT-small	BERT-small	TAPAS-small
	TAPAS-small	BERT-small	TAPAS-small

Table 2: Examined multi-modal bi-encoder and tri-encoder models.

	Bi-encoders	Tri-encoders
Learning rate	1e-5	1e-5
LR schedule	linear	linear
Warm-up steps	10%	10%
Batch size	38	28
Epochs	10	10
Optimizer	Adam	Adam

Table 3: Hyperparameters used to train the bi-encoder and tri-encoder multi-modal retrieval models.

the second architecture makes use of a BERT-small model for both the question encoder and the text encoder and utilizes a TAPAS-small instance to encode the tables. The third tri-encoder model uses two TAPAS-small instances to encode questions and tables and a BERT-small instance to encode text passages.

Herzig et al. (2021) use an additional down-projection layer to reduce the dimensionality of the question and table embeddings. We evaluated models with and without such a down-projection layer and found that their results do not differ significantly. Given that the models using an additional down-projection layer are more complex than models that directly utilize the embedding of the [CLS]-token, we consider only TAPAS-models without a down-projection layer throughout the remainder of this paper, including the experiments.

Table 3 specifies the hyperparameters used to train the bi-encoder and tri-encoder models on the training split of the MULTIMODALRETRIEVAL dataset described in Section 3. The learning objective is to create similar embeddings for relevant texts and/or tables with regard to a question. To train the models more efficiently, we make use of in-batch negatives besides each question’s hard neg-

ative text or table as suggested by Karpukhin et al. (2020) in the context of text retrieval. Given that the training samples inside a batch are randomly selected from all training examples, questions comprising a text passage as gold-label might have tables as negative labels, and vice versa.

5 Experiments

We compare the presented dense retrieval models to the sparse retrieval method BM25 and evaluate them based on recall@ k with $k \in \{10, 20, 100\}$. The search space to evaluate the models needs to consist of both texts and tables. For this purpose, 500,000 text passages are randomly sampled from Karpukhin et al. (2020)’s preprocessed Wikipedia passages making sure that the gold passages are among these passages. Furthermore, besides the text passages, all the tables from WIKISQL, OTT-QA, and NQ-TABLES are used, resulting in 656,166 tables and therefore approximately 1.2 million documents in total. The models are evaluated on a random sample of 1,000 questions of each dataset’s test split as listed in Table 1 and the full test split of the MULTIMODALRETRIEVAL dataset.

Following Karpukhin et al. (2020), a document retrieved for a question originating from NQ or NQ-TABLES is considered a correct match if the document contains the answer string of the granular answer. Given that the derivation of the granular answer for questions originating from WIKISQL and OTT-QA might need further aggregation, such as summation or counting, and, therefore, the answer string does not need to be present in a relevant document, a retrieved document is only considered a correct match if it is the gold annotated table. This evaluation procedure might have the effect of incorrectly judging non-gold tables that contain the answer to a query as irrelevant. However, since we apply the same evaluation procedure for all models, the numbers should be comparable. Table 4 specifies the evaluation results for BM25, all bi-encoder and all tri-encoder models.

The evaluation shows that BM25 outperforms all dense methods on both the full WIKISQL dataset and WIKISQL’s context-independent questions. It outperforms the best dense retrieval model on this dataset, the tri-encoder consisting of three BERT models, by 21.9 percentage points on all WIKISQL questions and 30.4 percentage points on context-independent WIKISQL questions with regard to

Encoders			NQ			WIKISQL			WIKISQL _{ctx-independent}		
Question	Text	Table	R@10	R@20	R@100	R@10	R@20	R@100	R@10	R@20	R@100
————	BM25	————	53.3	59.8	73.9	42.1	47.1	59.5	61.2	67.2	81.0
BERT	— BERT —		70.1	76.0	84.2	17.5	22.8	39.6	30.7	40.9	59.1
BERT	— TAPAS —		50.0	57.1	72.7	6.4	8.9	20.3	12.3	17.1	36.7
TAPAS	— TAPAS —		59.6	67.4	78.2	9.5	13.5	26.3	16.6	23.6	45.9
BERT	BERT	BERT	69.1	75.0	83.5	20.2	26.8	43.3	30.8	38.2	60.7
BERT	BERT	TAPAS	59.0	67.1	77.6	3.5	5.1	15.2	7.5	12.0	30.9
TAPAS	BERT	TAPAS	46.2	53.5	68.0	10.7	14.0	28.8	17.1	23.3	45.0

Encoders			OTT-QA			NQ-TABLES			MULTIMODALRETRIEVAL		
Question	Text	Table	R@10	R@20	R@100	R@10	R@20	R@100	R@10	R@20	R@100
————	BM25	————	40.2	45.6	58.2	56.6	65.1	82.8	50.7	57.0	71.1
BERT	— BERT —		72.9	78.0	89.4	84.9	91.2	96.7	55.2	61.8	73.8
BERT	— TAPAS —		30.6	40.0	63.1	71.6	72.6	90.0	34.2	39.1	56.6
TAPAS	— TAPAS —		49.9	60.4	82.8	78.9	86.0	94.0	42.9	50.2	65.4
BERT	BERT	BERT	73.8	79.7	90.1	86.4	91.6	96.7	56.1	62.3	74.9
BERT	BERT	TAPAS	25.8	34.3	59.4	52.3	62.1	80.7	29.6	36.1	52.8
TAPAS	BERT	TAPAS	50.8	60.9	79.8	76.9	82.8	92.9	40.3	46.9	62.9

Table 4: Evaluation results of BM25 and bi-encoder and tri-encoder retrieval models on 1000 random samples of the test splits of NQ, WIKISQL, context-independent questions of WIKISQL, OTT-QA, and NQ-TABLES and the full test set of our new MULTIMODALRETRIEVAL dataset with regard to recall@10, recall@20, and recall@100.

recall@10.

Analysing the WIKISQL dataset in more detail shows a very high lexical overlap of questions with their accompanying table. A combination of the Jaccard coefficient and Gestalt-Pattern-Matching² allows to quantify the lexical overlap of the questions with their corresponding tables without incorporating neither duplicate occurrences of the same word nor the order of the words inside the questions and the tables. This word order independence is particularly important for the analysis, given that the sparse retrieval method BM25 is order-agnostic. Even after lower-casing questions and tables and removing stopwords in the WIKISQL dataset, 40.68% of the questions lexically overlap completely with the relevant table, according to the combination of the Jaccard coefficient and Gestalt-Pattern-Matching. This large lexical overlap explains BM25’s strong performance on that dataset in Table 4. In contrast, only 0.27% of the questions in OTT-QA overlap completely with their accompanying table. For the other datasets, 15.73% of the questions in NQ-TABLES, 14.47% of the questions in NQ, and 21.96% of the ques-

²We use an implementation from: <https://github.com/seatgeek/fuzzywuzzy#token-set-ratio>

tions in MULTIMODALRETRIEVAL overlap completely with their accompanying table or text.

To better understand how the lexical overlap of questions and accompanying tables influences BM25’s performance, we split the test sets of WIKISQL and WIKISQL_{ctx-independent} into subsets with different ranges of lexical overlap. As can be observed in Figure 1, the recall of both BM25 and dense retrieval highly correlates with lexical overlap. Furthermore, while BM25 outperforms dense retrieval for questions with high lexical overlap, it is the other way round for questions with low lexical overlap.

On the sample of the NQ-TABLES test set, all dense retrieval models outperform BM25, except for the tri-encoder that consists of BERT instances as question and text passage encoder and a TAPAS instance as table encoder. The best performing model on this dataset is the tri-encoder consisting of three BERT encoders. This model outperforms the sparse retrieval method BM25 by 29.8 percentage points with regard to recall@10.

For the sampled questions of the OTT-QA development set, four out of the six dense retrieval models outperform BM25. The best performing model is again the tri-encoder model that is com-

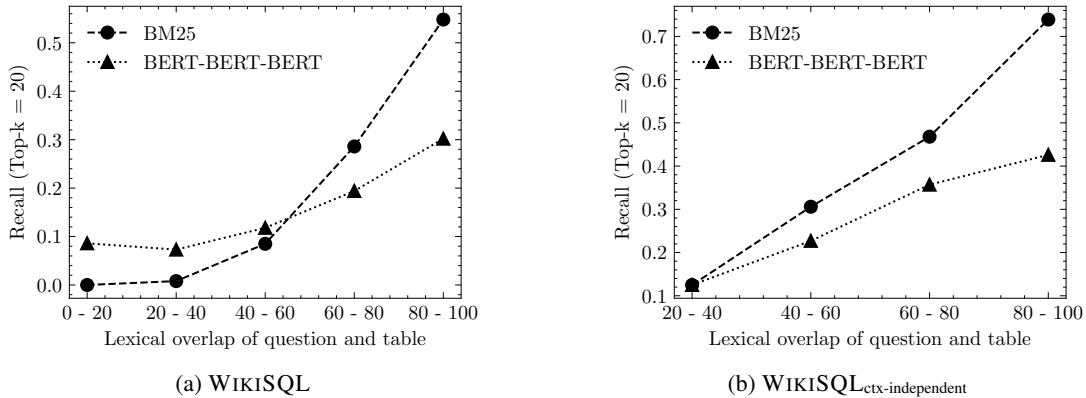


Figure 1: Recall@20 of the BM25 model and the BERT-BERT-BERT tri-encoder across different percentage of lexical overlap of question and table. Performance of BM25 drops drastically if the lexical overlap is low.

posed of three BERT-small encoders. This model outperforms BM25 by 33.6 percentage points with regard to recall@10. BM25 outperforms the bi-encoder consisting of a BERT model as question encoder and a TAPAS model as text and table encoder as well as the tri-encoder consisting of two BERT models serving as question encoder and text encoder, respectively, and a TAPAS model serving as table encoder.

When it comes to the performance on the text modality, i.e., questions deriving from NQ whose gold-label answer resides in a text passage, four out of the six dense retrieval models outperform BM25. For this case, the best performing model is not a tri-encoder but the bi-encoder comprising two BERT models. This model outperforms BM25 by 16.8 percentage points with regard to recall@10. However, this bi-encoder exceeds the tri-encoder consisting of three BERT encoders only slightly by one percentage point. The sparse retrieval method BM25 beats the bi-encoder consisting of a BERT model as question encoder and a TAPAS model as text and table encoder as well as the tri-encoder consisting of two TAPAS models serving as question encoder and table encoder, respectively, and a BERT model serving as text encoder.

In summary, the best performance on the WIKISQL test set is achieved by the sparse retrieval method BM25. The tri-encoder consisting of three BERT encoders shows the best performance on the remaining two tabular datasets, OTT-QA and NQ-TABLES. On the NQ dataset, i.e., questions whose answers reside in the textual modality, the bi-encoder consisting of two BERT encoders performs best but is almost on par with the tri-encoder consisting of three BERT models.

We can conclude that, under the limited experimental conditions, in particular, on the datasets used in our study, models involving TAPAS as question, text, and/or table encoder perform worse than models that rely only on BERT language models. Herzig et al. (2021) show that to be able to use TAPAS for the retrieval of tables, TAPAS needs to be additionally pre-trained on the table retrieval task. These pre-trained table retrieval models, which are used in the bi-encoders and tri-encoders that involve one or more TAPAS instances as encoder, are, however, pre-trained solely on the task of table retrieval and not text retrieval. Given the fact that the plain TAPAS model cannot be adapted to retrieval from scratch but needs this special pre-training, it might be the case that, to use TAPAS efficiently for the retrieval of both texts and tables, it needs to be pre-trained in a multi-modal setting on the retrieval of both texts and tables. Furthermore, batch size is significant for training retrieval models, as higher batch sizes make the training harder by adding more in-batch negatives. While the training of a bi-encoder does not allow a batch size higher than 38 and the training of a tri-encoder does not allow a batch size higher than 28 on a Tesla V100 GPU with 16 GB of memory, Herzig et al. (2021) make use of a batch size of 256 for training their TAPAS-based table retrieval models. Accordingly, it might be the case that TAPAS is more unstable to train and requires, therefore, larger batch sizes.

6 Conclusion and Future Work

This paper presented a transformer-based approach using bi-encoder and tri-encoder models for multi-modal retrieval of tables and texts. With experi-

ments on five datasets from related work and one newly created dataset, we show that the presented dense retrieval models outperform the sparse retrieval model BM25 if there is a low lexical overlap of questions and relevant tables and texts. More specifically, the tri-encoder architecture performs better on OTT-QA and NQ-TABLES, which represent the tabular modality, while the bi-encoder architecture performs slightly better on the NQ dataset representing the textual modality. We observe that the best retrieval models are those that rely only on BERT models as encoder and do not make use of TAPAS.

From an application point of view, future work could integrate the presented retrieval models as one component in a multi-modal open-domain QA pipeline and evaluate it on a real-world use case. Such a pipeline would facilitate information access immensely by combining valuable information from both sources rather than relying only on either texts or tables. Another promising path for future work is to extend our approach to more modalities with transformer-based models for images, videos, or speech. These models could serve as encoders for documents of different modalities to jointly train an n -encoder architecture, where one encoder is tailored to the queries and the remaining $n - 1$ encoders are tailored to each of the modalities that the user would like to search on. Last but not least, the research community would surely benefit from the creation of more multi-modal datasets to improve training and evaluation of multi-modal retrieval models and we are only making a first step in this direction with creating and releasing a dataset of tables and texts.

Acknowledgements

We would like to thank Jonathan Herzig and Julian Eisenschlos for taking the time to discuss ideas with us and to give early feedback on experiment results.

References

Ebrahim Bagheri and Feras Al-Obeidat. 2020. [A latent model for ad hoc table retrieval](#). In *Advances in Information Retrieval*, pages 86–93. Springer.

Michael J. Cafarella, Alon Halevy, Daisy Zhe Wang, Eugene Wu, and Yang Zhang. 2008. [Webtables: Exploring the power of tables on the web](#). *Proceedings of the VLDB Endowment*, 1(1):538–549.

Yinqiong Cai, Yixing Fan, Jiafeng Guo, Fei Sun, Ruqing Zhang, and Xueqi Cheng. 2021. [Semantic models for the first-stage retrieval: A comprehensive review](#). *arXiv preprint arXiv:2103.04831*.

Wenhu Chen, Ming-Wei Chang, Eva Schlinger, William Wang, and William Cohen. 2021. [Open question answering over tables and text](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Yang Wang. 2020a. [HybridQA: A dataset of multi-hop question answering over tabular and textual data](#). In *Findings of the Association for Computational Linguistics: EMNLP*, pages 1026–1036. Association for Computational Linguistics.

Zhiyu Chen, Mohamed Trabelsi, Jeff Heflin, Yinan Xu, and Brian D. Davison. 2020b. [Table search using a deep contextualized language model](#). In *Proceedings of the International Conference on Research and Development in Information Retrieval (SIGIR)*, page 589–598. Association for Computing Machinery.

Xiang Deng, Huan Sun, Alyssa Lees, You Wu, and Cong Yu. 2020. [Turl: Table understanding through representation learning](#). *Proceedings of the VLDB Endowment (PVLDB)*, 14(3):307–319.

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 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 4171–4186. Association for Computational Linguistics.

Julian Eisenschlos, Syrine Krichene, and Thomas Müller. 2020. [Understanding tables with intermediate pre-training](#). In *Findings of the Association for Computational Linguistics: EMNLP*, pages 281–296. Association for Computational Linguistics.

Jonathan Herzig, Thomas Müller, Syrine Krichene, and Julian Eisenschlos. 2021. [Open domain question answering over tables via dense retrieval](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 512–519. Association for Computational Linguistics.

Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. [TaPas: Weakly supervised table parsing via pre-training](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4320–4333. Association for Computational Linguistics.

- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics (ACL)*, 7:452–466.
- Alexander Hanbo Li, Patrick Ng, Peng Xu, Henghui Zhu, Zhiguo Wang, and Bing Xiang. 2021. [Dual reader-parser on hybrid textual and tabular evidence for open domain question answering](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 4078–4088. Association for Computational Linguistics.
- 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](#). *arXiv preprint arXiv:1907.11692*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, page 3111–3119. Curran Associates Inc.
- Feifei Pan, Mustafa Canim, Michael Glass, Alfio Gliozzo, and Peter Fox. 2021. [CLTR: An end-to-end, transformer-based system for cell-level table retrieval and table question answering](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, pages 202–209. Association for Computational Linguistics.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don't know: Unanswerable questions for SQuAD](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 784–789. Association for Computational Linguistics.
- Petar Ristoski and Heiko Paulheim. 2016. [RDF2Vec: RDF graph embeddings for data mining](#). In *Proceedings of the International Semantic Web Conference (ISWC)*, pages 498–514. Springer.
- Roei Shraga, Haggai Roitman, Guy Feigenblat, and Mustafa Canim. 2020a. [Ad hoc table retrieval using intrinsic and extrinsic similarities](#). In *Proceedings of The Web Conference (WWW)*, page 2479–2485. Association for Computing Machinery.
- Roei Shraga, Haggai Roitman, Guy Feigenblat, and Mustafa Canim. 2020b. [Web table retrieval using multimodal deep learning](#). In *Proceedings of the International Conference on Research and Development in Information Retrieval (SIGIR)*, page 1399–1408. Association for Computing Machinery.
- Alon Talmor, Ori Yoran, Amnon Catav, Dan Lahav, Yizhong Wang, Akari Asai, Gabriel Ilharco, Hananeh Hajishirzi, and Jonathan Berant. 2021. [Multimodalqa: Complex question answering over text, tables and images](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Ellen M Voorhees and Dawn M Tice. 2000. [The trec-8 question answering track report](#). In *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*.
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. [TaBERT: Pretraining for joint understanding of textual and tabular data](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 8413–8426. Association for Computational Linguistics.
- Li Zhang, Shuo Zhang, and Krisztian Balog. 2019. [Table2Vec: Neural word and entity embeddings for table population and retrieval](#). In *Proceedings of the International Conference on Research and Development in Information Retrieval (SIGIR)*, page 1029–1032. Association for Computing Machinery.
- Shuo Zhang and Krisztian Balog. 2018. [Ad hoc table retrieval using semantic similarity](#). In *Proceedings of the World Wide Web Conference (WWW)*, page 1553–1562. International World Wide Web Conferences Steering Committee.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. [Seq2SQL: Generating structured queries from natural language using reinforcement learning](#). *arXiv preprint arXiv:1709.00103*.