

Using the Hammer Only on Nails: A Hybrid Method for Representation-based Evidence Retrieval for Question Answering*

Zhengzhong Liang¹[0000-0003-0705-4880], Yiyun Zhao²[0000-0001-8280-8270], and
Mihai Surdeanu¹[0000-0001-6956-8030]

¹ Department of Computer Science, University of Arizona, Tucson AZ 85721, USA

² Department of Linguistics, University of Arizona, Tucson AZ 85719, USA
{zhengzhongliang, yiyunzhao, msurdeanu}@email.arizona.edu

Abstract. Evidence retrieval is a key component of explainable question answering (QA). We argue that, despite recent progress, transformer network-based approaches such as universal sentence encoder (USE-QA) do not always outperform traditional information retrieval (IR) methods such as BM25 for evidence retrieval for QA. We introduce a lexical probing task that validates this observation: we demonstrate that neural IR methods have the capacity to capture lexical differences between questions and answers, but miss obvious lexical overlap signal. Learning from this probing analysis, we introduce a hybrid approach for representation-based evidence retrieval that combines the advantages of both IR directions. Our approach uses a routing classifier that learns when to direct incoming questions to BM25 vs. USE-QA for evidence retrieval using very simple statistics, which can be efficiently extracted from the top candidate evidence sentences produced by a BM25 model. We demonstrate that this hybrid evidence retrieval generally performs better than either individual retrieval strategy on three QA datasets: OpenBookQA, ReQA SQuAD, and ReQA NQ. Furthermore, we show that the proposed routing strategy is considerably faster than neural methods, with a runtime that is up to 5 times faster than USE-QA.³

Keywords: Neural Information Retrieval, Representation-Based, BM25

1 Introduction

Open-domain question answering (QA) systems traditionally have three components: evidence retrieval, evidence reranking, and answer classification/extraction. In evidence retrieval, the model retrieves a smaller set of possibly useful evidence texts from a large knowledge base (KB), which are then reranked by the following component to push the most useful information to the top. Traditional directions

* This work was supported by the DARPA, grant number HR00111990011.

³ Code is available at: <https://github.com/clulab/releases/tree/master/ecir2021-hybrid-retrieval>

use word-overlap based models for evidence retrieval such as tf-idf and BM25. However, this can potentially cause the missing of useful information due to the “lexical chasm” [2] between the question and the answer. A potential remedy for this is to use neural networks for evidence retrieval, such as transformer network-based contextualized embedding methods [7, 28, 29].

Focusing on this evidence retrieval stage of a QA system, we argue that, for this component, transformer networks should not always be preferred over standard information retrieval (IR) methods. First, due to their reliance on continuous representations, transformer methods do not take direct advantage of obvious lexical evidence. This is a drawback in long-text retrieval, which tends to be affected less by the lexical chasm problem than short-text retrieval. Second, transformer-based methods are expensive to run, which makes them a less than ideal choice for end-user NLP applications with temporal constraints.

In this paper we introduce a *hybrid* approach for evidence retrieval for question answering. Our approach uses a routing classifier that routes an incoming question to either an IR method or a supervised transformer method for evidence retrieval, using solely shallow statistics sampled from the knowledge base of explanatory texts for each question. This strategy has two benefits: first, evidence retrieval performance improves overall because each question is handled by the appropriate retrieval method. Second, this method reduces computational overhead because for a considerable number of questions it does not use the more expensive neural component. In particular, our contributions are:

(1) We design and conduct a series of supervised lexical probing tasks on two QA datasets, which are trained to predict the terms in the query and the gold evidence text from the entire vocabulary, using as input either the tf-idf vector of the query, or the neural embedding of the same query. The comparison of the two probes indicates that the probe trained from the tf-idf vector of the query tends to predict terms that exist in the original query (thus emphasizing lexical overlap), whereas the probe trained on top of the query’s neural embedding predicts more terms in the evidence text that do not exist in the query (thus bridging the lexical chasm). This validates our hypothesis that different retrieval strategies should be used in different scenarios.

(2) Learning from this observation, we propose a hybrid retrieval method, which routes queries to either an information retrieval method (BM25 [24]) or a transformer-based one (USE-QA [28]). We show that this routing decision is learnable from simple statistics that can be efficiently extracted from the top documents retrieved by an IR method.

(3) We show that using this hybrid strategy generally improves evidence retrieval performance in three QA datasets: OpenBookQA [16], ReQA SQUAD, and ReQA Natural Questions (NQ) [1]. The hybrid approach performs significantly better than either individual model on ReQA SQUAD and NQ, with improvements in the mean reciprocal rank (MRR) of the correct evidence sentence ranging from 1% to 7.4% (depending on the dataset). On OpenBookQA the difference between the hybrid method and USE-QA is not statistically significant.

(4) Our analysis indicates that the hybrid method is significantly faster than neural IR methods. For example, the hybrid method is 2.2 times faster than USE-QA in OpenBookQA, and 5.2 times faster in ReQA SQUAD.

2 Related Work

Neural IR methods provide an exciting potential direction to mitigate the lexical chasm in QA [6]. Neural IR approaches can be broadly divided into two categories: representation-based and interaction-based [8]. Representation-based neural IR directions *pre-encode* the query and the document into a continuous representation learned using a subsample of the data, and use a shallow method to compute relevance scores at runtime (e.g., dot product) [12, 15]. Representation-based neural IR methods have low runtime overhead because all documents can be pre-computed as vectors, so that at test time the embeddings of the documents do not have to be recomputed for each query (i.e., the neural model is run for N_q times at test time, where N_q is the number of queries).

Interaction-based methods learn a query-specific representation of the documents *at runtime* [11, 19, 22]. Usually the query and the candidate document are concatenated and processed by a neural model jointly, so that complex interactions of the terms in the query and the document can be better captured. However, this requires running the neural model for $N_q \cdot N_d$ times at test time (where N_q is the number queries and N_d is the number of docs). Therefore interaction-based methods are not suitable for large-scale first stage retrieval and are usually used for second-stage retrieval (reranking). In this paper we focus on the representation-based method in the first stage retrieval.

Empirical evidence has shown that neural IR methods perform better in short-text retrieval, where the word-overlap-based IR methods are more likely to suffer from the lexical chasm problem. However, not much work has been done to show why neural IR methods are able to reduce the lexical chasm problem [8], partly because it is hard to explain the meaning of neural embeddings. Recently, probing tasks have been widely used to help understand the properties of neural networks [5, 9, 10]. In probing tasks, a shallow model is placed on top of the large neural model, and the shallow model is trained to show some properties of the large model. For example, in [10], the authors show that some syntactic information is encoded in the embeddings of the intermediate layers of BERT. Inspired by this, we design and conduct a series of lexical probing tasks to compare the abilities of traditional IR methods and neural IR methods to predict the terms that are indicative of lexical chasm, i.e., they exist in the evidence sentences but not in the original query.

Although they do not rely on explicit word overlaps, neural IR methods do not always outperform traditional IR. For instance, it has been shown that neural IR models usually work better on short text retrieval [4], and when training data is abundant [8], but not in other situations [13].

Efforts have been made to use traditional IR for evidence retrieval and neural IR for evidence reranking [3, 18, 21, 27]. However, *always* relying on traditional IR

Table 1. Statistics of the three datasets used throughout this paper, including the number of queries in the train/dev/test set, the number of candidate documents, and the average number of tokens per query/document.

Dataset	N train	N dev	N test	N doc	Avg. Q len.	Avg. D len.
OBQA	4,957	500	500	1,326	13.71	9.49
ReSQ	87,599	11,426	N/A	101,957	10.38	160.62
ReNQ	N/A	N/A	74,097	239,013	9.09	146.16

Table 2. Examples of queries, answer sentences, and contexts in the three datasets.

Dataset	Query	Answer Sentence	Context
ReSQ	To whom did the Virgin Mary allegedly appear in 1858 in Lourdes, France?	It is a replica of the grotto at Lourdes, France where the Virgin Mary reputedly appeared to Saint Bernadette Soubirous in 1858.	... a Marian place of prayer and reflection. It is a replica of the grotto at Lourdes, France where the Virgin Mary reputedly appeared to Saint Bernadette Soubirous in 1858. At the end of the main drive ...
ReNQ	Who sings the song "I don't care / I love it"?	In its chorus, Charli XCX shout in unison "I don't care / I love it".	... breaking up with an older boyfriend. In its chorus, Charli XCX shout in unison "I don't care / I love it". Critics compared the song's breakup ...
OBQA	Tadpoles start their lives as water animals	Tadpole changes into a frog	N/A

for evidence retrieval may miss useful evidence that does not have large lexical overlap with the query.

Motivated by these works, in this paper we propose a hybrid evidence retrieval direction for first-stage retrieval, in which we learn *when* to use traditional IR vs. neural IR. As our results show, this yields a more accurate retrieval component that also has a lower runtime overhead than neural methods.

3 Datasets and Evaluation Measures

We conduct our probing analyses and retrieval experiments on three QA-related retrieval datasets. One of these datasets comes from the science domain; the other two are open domain. More statistics and examples of these datasets are shown in table 1 and 2, and we describe them below.

OpenBookQA: The OpenBookQA dataset [16] (abbreviated to *OBQA* from now on) addresses a multiple-choice QA task in the science domain. Each correct answer is jointly annotated with one key evidence sentence (or *justification*) that supports its correctness. The justification comes from a knowledge base of 1326 sentences. In this paper, we construct a corpus of 1326 documents from these sentences. Further, for each question, we concatenate the question and the correct answer choice to form the *query*, and retrieve the gold justification (or *target document*) for that query from the corpus of 1326 documents.

ReQA SQuAD: The ReQA SQuAD dataset [1] (abbreviated to *ReSQ*) is a sentence-level retrieval dataset converted from the SQuAD reading comprehension dataset [23]. In the original SQuAD reading comprehension task, the answers to questions must be extracted from sentences in a set of provided paragraphs. The ReQA SQuAD dataset uses the questions in SQuAD as the queries, and converts all paragraphs to single sentences. The goal of this retrieval task is to retrieve

Table 3. Results of the probing tasks on two datasets. We report mean average precision (MAP) (higher is better) and perplexity (PPL) (lower is better) scores for the gold terms to be predicted. We report separate scores for terms in the query, and terms that occur *only* in the justification fact (mean and stdev across 5 random seeds).

Dataset Task		Query MAP	Query PPL	Fact MAP	Fact PPL
OBQA	USE-QA embd, gold label	0.306 ± 0.01	1.709 ± 0.02	0.154 ± 0.01	1.188 ± 0.01
	tf-idf embd, gold label	0.458 ± 0.01	1.558 ± 0.01	0.098 ± 0.00	1.334 ± 0.01
	Random embd, gold label	0.053 ± 0.02	3.640 ± 0.10	0.031 ± 0.01	3.294 ± 1.63
	USE-QA embd, rand label	0.085 ± 0.00	1.974 ± 0.02	0.043 ± 0.00	1.557 ± 0.02
ReSQ	USE-QA embd, gold label	0.139 ± 0.01	1.944 ± 0.01	0.147 ± 0.00	2.134 ± 0.00
	tf-idf embd, gold label	0.142 ± 0.04	1.828 ± 0.00	0.127 ± 0.01	2.043 ± 0.00
	Random embd, gold label	0.091 ± 0.00	7.419 ± 3.50	0.093 ± 0.00	3.478 ± 0.15
	USE-QA embd, rand label	0.122 ± 0.02	1.909 ± 0.00	0.124 ± 0.00	2.013 ± 0.00

the sentence that contains the correct answer from all the sentences generated from all the paragraphs. Since some answer sentences are meaningless without the surrounding context, each candidate sentence is accompanied by its original paragraph as the context.

ReQA NQ: The ReQA NQ dataset [1] (abbreviated to *ReNQ*) is similarly converted from another reading comprehension task – Natural Questions [14] – following the same process as ReQA SQuAD. Similarly, each query is a question and each target document is a sentence/context pair, where the context is the paragraph that contains the gold justification.

4 Understanding the Behavior of Neural IR through Lexical Probing

Our key hypothesis is that neural IR methods are better at modeling the lexical chasm between queries and evidence sentences than traditional IR, whereas traditional IR captures explicit lexical overlap better. We design a lexical probe and control tasks to investigate this.

4.1 Task Overview

Figure 1 summarizes our lexical probe, with an example from OBQA.

Probe input: The probe starts by generating a representation of the input query. This representation is either: (a) the tf-idf vector of the query, generated using *scikit-learn* [20], or (b) the query embedding generated by USE-QA [28].

Linear layer: This vector is fed to a linear layer, with input size N_d and output size N_v , to predict the terms (i.e., unique words) in the query *and* in the gold fact. N_v is the size of vocabulary \mathcal{V} , where \mathcal{V} is the set of all terms in the dataset. Each

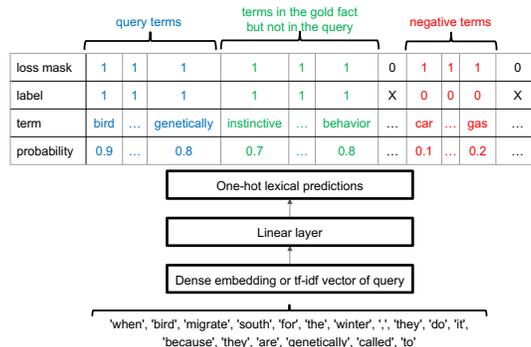


Fig. 1. Probe task overview: the linear probe is trained to predict the terms in the query and in the gold fact from the entire vocabulary, given either the input embedding or the tf-idf vector of the query. This probe investigates the capability of the query representation to predict both lexical overlap (i.e., terms from the query), as well as its ability to bridge the lexical chasm between queries and supporting facts (i.e., predict terms that exist in the fact and not in the query). A loss mask is used to make sure the loss is only computed on certain terms during training.

number in the output is the predicted probability of that particular term being in the query/gold fact. Note that the input embedding/vector is not changed during the training of the probe task. Thus, if the neural embedding contains meaningful information about the gold fact, it should perform better than tf-idf on predicting the terms that are only in the gold fact.

Training label and loss: For each query q_i , we use \mathcal{P}_i to indicate the set of all terms in $\{q_i, \text{gold_fact}(q_i)\}$. The training label is a one-hot vector of size N_v , where the values for the terms in \mathcal{P}_i are 1, and the rest of the entries are 0. However, since the terms in \mathcal{P}_i are considerably fewer than the whole vocabulary, there will be many more 0s than 1s in this label vector, causing label imbalance. Therefore, we construct a set of negative terms \mathcal{N}_i , which contains terms that are randomly sampled from the vocabulary \mathcal{V} but not in \mathcal{P}_i . The size of \mathcal{N}_i also equals to the size of \mathcal{P}_i . The loss is only computed on the terms in $\mathcal{P}_i \cup \mathcal{N}_i$ instead of the whole vocabulary \mathcal{V} . The total loss of each query q_i is summarized by: $L = -\sum_{j \in \mathcal{P}_i \cup \mathcal{N}_i} [y_j \log \hat{y}_j + (1 - y_j) \log(1 - \hat{y}_j)]$, where $y_j \in \{0, 1\}$ is the label, and $\hat{y}_j \in (0, 1)$ is the predicted probability of the corresponding term.

4.2 Control Tasks

We designed two control tasks to check whether the information necessary for prediction is contained in the query representation and not in the linear layer [9]:

Random embedding (rand embd): This probe replaces the neural embedding with a randomly-generated embedding. If the query representation encodes useful

information, this probe should perform much worse than the one using the neural representation.

Rand label (rand label): In this experiment we randomly replace the target terms in both training and testing. For example, we replace the terms to be predicted for query i with terms from a randomly-selected query j : $\{q_j, \text{gold_fact}(q_j)\}$. This is to examine whether it is possible for the linear probe to learn non-sensical associations between random (embedding, target terms) pairs.

4.3 Probe Results

Table 3 lists the results of these probing tasks. We draw several observations:

(1) With minor exceptions, the two actual probes perform better than the two control tasks. This confirms that there is indeed signal that is encoded in the query representations, and this is what the linear probe classifier exploits.

(2) The probe that relies on the neural query representation obtains higher fact MAP (and lower fact PPL) than the probe that uses the tf-idf representation. This indicates that the neural representation does indeed contain information that helps bridge queries, answers, and supporting facts. On the other hand, the tf-idf probe has higher query MAP (and lower query PPL) than the neural probe. This confirms that the traditional IR representation is better at capturing explicit lexical overlap with the query than the neural one. All in all, these observations suggest that these two retrieval directions are better at different things.

5 Hybrid Retrieval Approach

5.1 Individual IR Models

The hybrid approach proposed builds from (and is compared against) the following individual retrieval models. Note that these approaches were chosen because they had the best performance on these datasets. For example, USE-QA consistently performed better than BERT.

BM25: We use the Lucene 6.4.0⁴ Java implementation of BM25 [24] as the “traditional” IR method. For OBQA, each document is one sentence in the knowledge base corpus (1326 sentences in total). In ReSQ and ReNQ, each document is constructed by concatenating the candidate answer sentence and its context (so that each candidate answer sentence appears twice in the document).

BERT: For this method we fine-tune a pretrained BERT-base model [7, 26]. We use the BERT retriever in the representation-based manner: the query q_i and the document d_j are encoded using the [CLS] embedding of BERT as h_i^q and h_j^d . Then the relevance score of q_i and d_j are obtained by $Rel(q_i, f_j) = h_i^q \cdot h_j^d$. For

⁴ <https://lucene.apache.org>

ReSQ and ReNQ, each document is composed of the candidate answer sentence and its context. We concatenate them and separate them with the [SEP] token. Therefore, the input of each document is “[CLS] candidate answer sentence [SEP] context sentences [SEP]”.

USE-QA: The USE-QA retriever [28] has separate encoders for the query and document. The query encoder is a transformer-based model, producing a 512-dimension embedding as the query representation. The document encoder has a transformer-based model to encode the answer sentence and a Convolutional Neural Network (CNN)-based model to encode the context. A single 512-dimension embedding is produced as the document representation. Finally, the relevance score is computed as the dot product of the query embedding and the document embedding. USE-QA is pre-trained on large scale retrieval tasks and, as used in [1], we do not fine-tune it in the retrieval tasks.

5.2 Are Neural IR Methods Generally Better than Traditional IR?

The probe introduced in Section 4 indicates that neural and traditional IR methods have different behaviors. But what impact does that have in practice, with respect to overall performance? To answer this question, we performed a comparison that aims to understand if transformer-based retrieval methods are better overall than traditional IR. Due to space limitations, we discuss here results from the best individual models in each class: BM25 for traditional IR, and USE-QA for neural IR (We observed similar behavior from tf-idf and BERT.). We use two datasets: one domain-specific (OBQA) and one open-domain (ReSQ).

Figure 2 summarizes this comparison between BM25 and USE-QA on the dev partitions of OBQA and on a 10,000-query subset of ReSQ training partition. Here, we consider a model better than the other when it yields a better ranking for the correct justification. We draw two observations from this analysis:

(1) No approach is consistently better than the other. Overall, BM25 is at least as good as USE-QA in 293 queries out of 500 queries in OBQA dev set, and 7603 queries out of 10000 randomly sampled queries in ReSQ train set. This is further motivation for a hybrid approach.

(2) Importantly, there is immediate signal to differentiate between the two situations. When BM25 performs better than or similarly to USE-QA, the top BM25 score (after the softmax normalization) tends to be in the 0.8 to 1 range. In contrast, when there is little lexical overlap between question and justification indicated by low BM25 scores, e.g., below 0.2, USE-QA performs considerably better. This supports the intuition that USE-QA can capture lexical differences between question and justification when present.

5.3 Hybrid Retrieval Model

Motivated by the previous observations, we propose a hybrid evidence retrieval method that uses a routing classifier to direct an incoming question to either

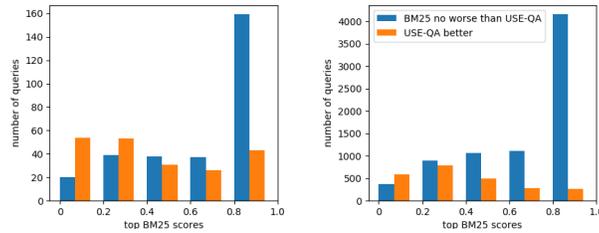


Fig. 2. Histograms of queries in OBQA dev (left) and a randomly sampled subset of ReSQ train (right) where BM25 is no worse than USE-QA (blue) or where USE-QA is better (orange). The x axis is the top BM25 score after a softmax is applied to the BM25 scores of the top 64 sentences.

the BM25 retriever or the USE-QA retriever based on simple statistics that can be extracted efficiently. The key intuition behind our hybrid strategy is that we can estimate the optimal retrieval method based on the top answers retrieved by traditional IR. In particular, if these answers receive a high score from the traditional IR method, it indicates that the current scenario is driven by lexical overlap, and traditional IR is likely to do better; the opposite is true otherwise.

We propose two variants for the routing classifier:

Hybrid (1-param): This classifier uses a single parameter: a threshold on the normalized BM25 score of top document retrieved by the BM25 method.⁵ If the top 1 score is higher than this threshold, the classifier routes the question to BM25; otherwise it sends it to USE-QA. This is a simple implementation of the intuition above – if the top normalized BM25 score is high, then it is likely that there is a candidate document that has a large lexical overlap with the query, and which is probably a correct justification. On the other hand, if the top normalized BM25 score is low, it is either because: (a) there is no document that has a large lexical overlap with the query, or (b) because there are multiple candidate documents that have high BM25 scores (and they are squished during normalization). In either of these scenarios BM25 is unlikely to identify the gold document, and, therefore, USE-QA should be selected. The value of this threshold is determined by performing a grid search on the dev partition, with the threshold ranging from 0 to 1 with an interval of 0.1.

Hybrid (BM25): This classifier is a generalization of the above. That is, instead of relying solely on the top retrieved document, this classifier extracts features from the top k . In particular, for each query, we construct a feature vector f and use a logistic regression (LR) classifier that takes f to predict whether to use BM25 or USE-QA. The i^{th} feature of f is computed as $f_i = mean(S[0 : 2^i])$, where S are the top BM25 scores ranked in the descending order (after softmax normalization). In this paper we use i up to 6 (i.e., use up to top 64 BM25

⁵ We normalize this score by applying a softmax layer to the BM25 scores of the top k ($k = 64$ in this paper) documents.

scores). For example, feature 2 averages the BM25 scores of the top 4 documents retrieved by the traditional IR method. This strategy allows the classifier to take advantage of more documents when needed, but also focus on the top result(s) when they are sufficiently predictive.

Note that all *Hybrid* approaches choose *either* one of the individual models. USE-QA is *not* used as a reranking method on top of BM25, because USE-QA is applied to all documents instead of the top documents retrieved by BM25.

6 Results

In this section, we empirically evaluate the proposed evidence retrieval methods. We use the mean reciprocal rank (MRR) score [25] of the correct evidence sentence (or target document) in the test dataset as the evaluation measure.

Since ReSQ only provides training and developments partitions, we randomly sample 10,000 queries from the training data and use them for development, and use the original development set of ReSQ as test. ReNQ does not provide training/development/test partitions; for this dataset we use 5-fold cross-validation, sampling 10,000 queries from one fold as the development data in each split, and using the remaining folds as test.

For all datasets, USE-QA is used without fine-tuning as proposed in [1]. For the BERT retriever, we fine-tune it on the training data of OBQA and ReSQ. For all hybrid retrievers, we tune their routing classifiers on the respective development partitions. For ReSQ, we further divide the development set into 5 splits (2,000 queries in each split) and tune 5 routing classifiers and evaluate them separately on the full test set to make sure the results are robust to different development sets.

6.1 Individual vs Hybrid Retrievers

Table 4 shows the MRR scores of the individual retrieval methods compared to the hybrid ones, on the three datasets. We draw several observations from this:

(1) Most hybrid strategies outperform the individual retrieval methods, as well as the naive strategy that simply sums up the scores of two individual models, and uses the sum for ranking. *Hybrid (1-param)* and *Hybrid (BM25)* are statistically significantly better than BM25 and USE-QA on ReSQ and ReNQ under a bootstrap resampling significance analysis (10,000 samples, p -value $< 10^{-5}$). On OBQA, *Hybrid (1-param)* and *Hybrid (BM25)* are statistically significantly better than BM25 under the same bootstrap resampling significance analysis, but there is no significant difference between the hybrid methods and USE-QA. This demonstrates that transformer-based and IR-based methods capture complementary information, and the distinction of when to use one vs. another is learnable. Table 5 lists several runtime statistics of our best classifier, *Hybrid (BM25)*, which support this observation. The first two rows indicate that the routing classifier uses both individual retrievers, with around 60% (OBQA) or

Table 4. Mean reciprocal rank (MRR) scores of the retrieval methods investigated on the three QA datasets. The *BM25 + USE-QA* method sums up the scores produced by BM25 and USE-QA, and uses that score for ranking. * indicates that *Hybrid (BM25)* is statistically significantly better than both USE-QA and BM25 (bootstrap resampling with 10,000 iterations; p -value $< 10^{-5}$). *Hybrid-NN (BM25)* uses approximate Nearest Neighbor for USE-QA in the hybrid method with 20 search trees. The *Ceiling* method always selects the best individual model (USE-QA or BM25) for each query by their ranking of the gold justification.

	BM25	BERT	USE-QA	BM25 + USE-QA	Hybrid (1-param)	Hybrid (BM25)	Hybrid-NN (BM25)	Ceiling
OBQA	0.522	0.557	0.610	0.550	0.611	0.596	N/A	0.69
ReSQ	0.645	0.260	0.520	0.647	0.656	0.657*	0.656*	0.71
ReNQ	0.293	N/A	0.223	0.290	0.301	0.303*	0.298*	0.39

Table 5. Routing statistics for the routing classifier that trains a logistic regression model using features extracted from the top 64 BM25 documents.

	OBQA	ReSQ	ReNQ
n samples routed to BM25	306	49,270	260,640
n samples routed to useQA	194	7,860	59,845
Samples w/ improved rankings vs. BM25	129	4,095	33,078
Samples w/ worse rankings vs. BM25	47	3,257	25,562
Samples w/ improved rankings vs. useQA	61	21,488	146,020
Samples w/ worse rankings vs. useQA	93	8,640	84,839

86% (ReSQ) of questions being routed to BM25. The next four rows indicate that, on average, the hybrid approach improves over both individual methods especially on ReSQ and ReNQ.

(2) While *Hybrid (BM25)* outperforms the simpler *Hybrid (1-param)*, the difference is not statistically significant.⁶ This further suggests that simpler approaches work in this case. The routing decision can be approximated with a single parameter (a threshold on the BM25 score), applied to a single justification that is efficiently extracted by IR.

6.2 Runtime Analysis

A further advantage of our hybrid approach is improved runtime over neural methods, because a considerable number of queries are routed to a traditional, fast IR engine. To investigate this, we measure the processing time per query using BM25, USE-QA and various hybrid retrievers and calculate the total time usage of these retrieval methods. The processing time per query is measured as:

(1) For BM25, we measure the time of parsing the query, searching the top k ($k = 1400$ for OBQA, and 2000 for ReSQ and ReNQ) documents, and sorting the retrieved documents by the BM25 scores.

⁶ Bootstrap resampling with 10,000 samples, p -value < 0.13

Table 6. Runtime comparison of BM25, USE-QA and hybrid retrievers on the corresponding test partitions. All times are the total times in seconds on all test queries.

	BM25	USE-QA	BM25 + USE-QA	Hybrid (1-param)	Hybrid (BM25)	Hybrid-NN (BM25)
OBQA	1.38	19.74	21.23	20.85	8.95	N/A
ReSQ	179.56	3241.73	3476.99	922.28	625.74	593.82
ReNQ	1547.56	26722.63	28929.47	9513.97	6565.21	3696.11

(2) For USE-QA, we measure the query processing time (including query tokenization and the embedding generation of USE-QA⁷), searching the top k (1326 for OBQA and 2000 for ReSQ and ReNQ) documents, and sorting them by the scores. We run this experiment using Tensorflow on Google Colab with GPU.

(3) For hybrid models, the processing time of each query is the sum of: (1) the BM25 processing time (2) the runtime of the routing classifier and (3) the processing time of USE-QA if USE-QA is selected for that query.

Table 6 shows the total processing time of all queries using different retrieval methods. The table indicates that USE-QA is more than 15 times slower than BM25 on all datasets. Further, the hybrid approach reduces that gap while still allowing for the benefits of the neural IR when needed: *Hybrid (BM25)* is 2.2 times faster than USE-QA in OBQA, and 5.2 times faster in ReSQ. Our hybrid method is also significantly faster than *BM25 + USE-QA*, which uses both BM25 and the neural retriever on every query [17].

7 Conclusion

We argue that transformer network-based approaches do not always outperform IR methods for evidence retrieval for QA. We validate this observation with an empirical analysis, and with a lexical probing task where two probes were trained to predict words in the gold evidence texts. The first probe, trained on the tf-idf vector of the query, tends to predict words that exist in the original query (thus emphasizing lexical overlap), whereas the second probe, trained on top of the query’s neural embedding, predicts more words in the evidence text that do not exist in the query (bridging the lexical differences between these texts).

Learning from this analysis, we introduced a routing classifier that learns when to direct incoming questions to traditional or neural IR methods for evidence retrieval. The routing classifier is trained using very simple statistics, which can be extracted from the top candidate evidence sentences produced by traditional IR. We showed that this hybrid evidence retrieval generally performs better than either individual retrieval strategy on three QA datasets. Further, we showed that this routing classifier can be approximated with nearly the same performance with a 1-parameter model (a threshold over the IR score of the top evidence sentence retrieved by BM25), which simplifies real-world applications of our approach. Lastly, we show that our routing classifier is considerably faster than USE-QA, with runtime improvements of up to 5 times.

⁷ The batch size is set to 1 when generating the embedding, for a fair comparison with BM25, and because in a real use case the queries may not arrive in batch.

References

1. Ahmad, A., Constant, N., Yang, Y., Cer, D.: Reqa: An evaluation for end-to-end answer retrieval models. In: Proceedings of the 2nd Workshop on Machine Reading for Question Answering. pp. 137–146 (2019)
2. Berger, A., Caruana, R., Cohn, D., Freitag, D., Mittal, V.: Bridging the lexical chasm: statistical approaches to answer-finding. In: Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval. pp. 192–199. ACM (2000)
3. Chen, R.C., Spina, D., Croft, W.B., Sanderson, M., Scholer, F.: Harnessing semantics for answer sentence retrieval. In: Proceedings of the Eighth Workshop on Exploiting Semantic Annotations in Information Retrieval. pp. 21–27. ACM (2015)
4. Cohen, D., Ai, Q., Croft, W.B.: Adaptability of neural networks on varying granularity in tasks. arXiv preprint arXiv:1606.07565 (2016)
5. Conneau, A., Kruszewski, G., Lample, G., Barrault, L., Baroni, M.: What you can cram into a single $\&!#$ vector: Probing sentence embeddings for linguistic properties. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2126–2136 (2018)
6. Dehghani, M., Zamani, H., Severyn, A., Kamps, J., Croft, W.B.: Neural ranking models with weak supervision. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 65–74. ACM (2017)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: 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). pp. 4171–4186 (2019)
8. Guo, J., Fan, Y., Pang, L., Yang, L., Ai, Q., Zamani, H., Wu, C., Croft, W.B., Cheng, X.: A deep look into neural ranking models for information retrieval. *Information Processing & Management* **57**(6), 102067 (2020)
9. Hewitt, J., Liang, P.: Designing and interpreting probes with control tasks. 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). pp. 2733–2743 (2019)
10. Hewitt, J., Manning, C.D.: A structural probe for finding syntax in word representations. 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). pp. 4129–4138 (2019)
11. Hu, B., Lu, Z., Li, H., Chen, Q.: Convolutional neural network architectures for matching natural language sentences. In: Advances in neural information processing systems. pp. 2042–2050 (2014)
12. Huang, P.S., He, X., Gao, J., Deng, L., Acero, A., Heck, L.: Learning deep structured semantic models for web search using clickthrough data. In: Proceedings of the 22nd ACM international conference on Information & Knowledge Management. pp. 2333–2338. ACM (2013)
13. Iyyer, M., Manjunatha, V., Boyd-Graber, J., Daumé III, H.: Deep unordered composition rivals syntactic methods for text classification. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). vol. 1, pp. 1681–1691 (2015)

14. Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., et al.: Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics* **7**, 453–466 (2019)
15. Lee, K., Chang, M.W., Toutanova, K.: Latent retrieval for weakly supervised open domain question answering. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. pp. 6086–6096 (2019)
16. Mihaylov, T., Clark, P., Khot, T., Sabharwal, A.: Can a suit of armor conduct electricity? a new dataset for open book question answering. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. pp. 2381–2391 (2018)
17. Mitra, B., Diaz, F., Craswell, N.: Learning to match using local and distributed representations of text for web search. In: *Proceedings of the 26th International Conference on World Wide Web*. pp. 1291–1299 (2017)
18. Nie, Y., Wang, S., Bansal, M.: Revealing the importance of semantic retrieval for machine reading at scale. 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)*. pp. 2553–2566 (2019)
19. Nogueira, R., Cho, K.: Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085* (2019)
20. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
21. Pirttoaca, G.S., Rebedea, T., Ruseti, S.: Answering questions by learning to rank-learning to rank by answering questions. 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)*. pp. 2531–2540 (2019)
22. Qiao, Y., Xiong, C., Liu, Z., Liu, Z.: Understanding the behaviors of bert in ranking. *arXiv preprint arXiv:1904.07531* (2019)
23. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250* (2016)
24. Robertson, S., Zaragoza, H.: The probabilistic relevance framework: Bm25 and beyond. *Found. Trends Inf. Retr.* **3**(4), 333–389 (Apr 2009). <https://doi.org/10.1561/1500000019>, <https://doi.org/10.1561/1500000019>
25. Voorhees, E.: *Proceedings of the 8th text retrieval conference. TREC-8 Question Answering Track Report* pp. 77–82 (1999)
26. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Brew, J.: Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv abs/1910.03771* (2019)
27. Yang, W., Xie, Y., Lin, A., Li, X., Tan, L., Xiong, K., Li, M., Lin, J.: End-to-end open-domain question answering with bertserini. *NAACL HLT 2019* p. 72 (2019)
28. Yang, Y., Cer, D., Ahmad, A., Guo, M., Law, J., Constant, N., Abrego, G.H., Yuan, S., Tar, C., Sung, Y.H., et al.: Multilingual universal sentence encoder for semantic retrieval. *arXiv preprint arXiv:1907.04307* (2019)
29. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: Xlnet: Generalized autoregressive pretraining for language understanding. In: *Advances in neural information processing systems*. pp. 5753–5763 (2019)