

Joint Models for Answer Verification in Question Answering Systems

Zeyu Zhang*, Thuy Vu, and Alessandro Moschitti

School of Information, The University of Arizona, Tucson, AZ, USA

Amazon Alexa AI, Manhattan Beach, CA, USA

zeyuzhang@email.arizona.edu, {thuyvu, amosch}@amazon.com

Abstract

This paper studies joint models for selecting correct answer sentences among the top k provided by answer sentence selection (AS2) modules, which are core components of retrieval-based Question Answering (QA) systems. Our work shows that a critical step to effectively exploit an answer set regards modeling the interrelated information between pair of answers. For this purpose, we build a three-way multi-classifier, which decides if an answer supports, refutes, or is neutral with respect to another one. More specifically, our neural architecture integrates a state-of-the-art AS2 model with the multi-classifier, and a joint layer connecting all components. We tested our models on WikiQA, TREC-QA, and a real-world dataset. The results show that our models obtain the new state of the art in AS2.

1 Introduction

Automated Question Answering (QA) research has received a renewed attention thanks to the diffusion of Virtual Assistants. Among the different types of methods to implement QA systems, we focus on Answer Sentence Selection (AS2) research, originated from TREC-QA track (Voorhees and Tice, 1999), as it proposes efficient models that are more suitable for a production setting, e.g., they are more efficient than those developed in machine reading (MR) work (Chen et al., 2017).

Garg et al. (2020) proposed the TANDA approach based on pre-trained Transformer models, obtaining impressive improvement over the state of the art for AS2, measured on the two most used datasets, WikiQA (Yang et al., 2015) and TREC-QA (Wang et al., 2007). However, TANDA was applied only to pointwise rerankers (PR), e.g., simple binary classifiers. Bonadiman and Moschitti

*Work done while the author was an intern at Amazon Alexa

<i>Claim:</i>	Joe Walsh was inducted in 2001.
<i>Ev</i> ₁ :	As a member of the Eagles, Walsh was inducted into the Rock and Roll Hall of Fame in 1998, and into the Vocal Group Hall of Fame in 2001.
<i>Ev</i> ₂ :	Joseph Fidler Walsh (born November 20, 1947) is an American singer songwriter, composer, multi-instrumentalist and record producer.
<i>Ev</i> ₃ :	Walsh was awarded with the Vocal Group Hall of Fame in 2001.

Table 1: A claim verification example from FEVER.

(2020) tried to improve this model by jointly modeling all answer candidates with listwise methods, e.g., (Bian et al., 2017). Unfortunately, merging the embeddings from all candidates with standard approaches, e.g., CNN or LSTM, did not improve over TANDA.

A more structured approach to building joint models over sentences can instead be observed in Fact Verification Systems, e.g., the methods developed in the FEVER challenge (Thorne et al., 2018a). Such systems take a claim, e.g., *Joe Walsh was inducted in 2001*, as input (see Tab. 1), and verify if it is valid, using related sentences called *evidences* (typically retrieved by a search engine). For example, *Ev*₁, *As a member of the Eagles, Walsh was inducted into the Rock and Roll Hall of Fame in 1998, and into the Vocal Group Hall of Fame in 2001*, and *Ev*₃, *Walsh was awarded with the Vocal Group Hall of Fame in 2001*, support the veracity of the claim. In contrast, *Ev*₂ is neutral as it describes who Joe Walsh is but does not contribute to establish the induction. We conjecture that supporting evidence for answer correctness in AS2 task can be modeled with a similar rationale.

In this paper, we design joint models for AS2 based on the assumption that, given q and a target answer candidate t , the other answer candidates, (c_1, \dots, c_k) can provide positive, negative, or neutral support to decide the correctness of t . Our first approach exploits Fact Checking research: we adapted a state-of-the-art FEVER system, KGAT (Liu et al., 2020), for AS2. We defined a claim as

a pair constituted of the question and one target answer, while considering all the other answers as *evidences*. We re-trained and rebuilt all its embeddings for the AS2 task.

Our second method, Answer Support-based Reranker (ASR), is completely new, it is based on the representation of the pair, (q, t) , generated by state-of-the-art AS2 models, concatenated with the representation of all the pairs (t, c_i) . The latter summarizes the contribution of each c_i to t using a max-pooling operation. c_i can be unrelated to (q, t) since the candidates are automatically retrieved, thus it may introduce just noise. To mitigate this problem, we use an Answer Support Classifier (ASC) to learn the relatedness between t and c_i by classifying their embedding, which we obtain by applying a transformer network to their concatenated text. ASC tunes the (t, c_i) embedding parameters according to the evidence that c_i provides to t . Our Answer Support-based Reranker (ASR) significantly improves the state of the art, and is also simpler than our approach based on KGAT.

Our third method is an extension of ASR. It should be noted that, although ASR exploits the information from the k candidates, it still produces a score for a target t without knowing the scores produced for the other target answers. Thus, we jointly model the representation obtained for each target in a multi-ASR (MASR) architecture, which can then carry out a complete global reasoning over all target answers.

We experimented with our models over three datasets, WikiQA, TREC-QA, and Web-based Question Answering (WQA), where the latter is an internal dataset built on anonymized customer questions¹. The results show that:

- ASR improves the best current model for AS2, i.e., TANDA by $\sim 3\%$, corresponding to an error reduction of 10% in Accuracy, on both WikiQA and TREC-QA.
- We also obtain a relative improvement of $\sim 3\%$ over TANDA on WQA, confirming that ASR is a general solution to design accurate QA systems.
- Most interestingly, MASR improves ASR by additional 2%, confirming the benefit of joint modeling.

¹We are making it available soon

Finally, it is interesting to mention that MASR improvement is also due to the use of FEVER data for pre-fine-tuning ASC, suggesting that the fact verification inference and the answer support inference are similar.

2 Problem definition and related work

We consider retrieval-based QA systems, which are mainly constituted by (i) a search engine, retrieving documents related to the questions; and (ii) an AS2 model, which reranks passages/sentences extracted from the documents. The top sentence is typically used as final answer for the users.

2.1 Answer Sentence Selection (AS2)

The task of reranking answer-sentence candidates provided by a retrieval engine can be modeled with a classifier scoring the candidates. Let q be an element of the question set, Q , and $\mathcal{A} = \{c_1, \dots, c_n\}$ be a set of candidates for q , a reranker can be defined as $\mathcal{R} : Q \times \Pi(\mathcal{A}) \rightarrow \Pi(\mathcal{A})$, where $\Pi(\mathcal{A})$ is the set of all permutations of \mathcal{A} . Previous work targeting ranking problems in the text domain has classified reranking functions into three buckets: pointwise, pairwise, and listwise methods.

Pointwise reranking: This approach learns $p(q, c_i)$, which is the probability of c_i correctly answering q , using a standard binary classification setting. The final rank is simply obtained sorting c_i , based on $p(q, c_i)$. Previous work estimates $p(q, c_i)$ with neural models (Severyn and Moschitti, 2015), also using attention mechanisms, e.g., Compare-Aggregate (Yoon et al., 2019), inter-weighted alignment networks (Shen et al., 2017), and pre-trained Transformer models, which are the state of the art. Garg et al. (2020) proposed TANDA, which is the current most accurate model on WikiQA and TREC-QA.

Pairwise reranking: The method considers binary classifiers of the form $\chi(q, c_i, c_j)$ for determining the partial rank between c_i and c_j , then the scoring function $p(q, c_i)$ is obtained by summing up all the contributions with respect to the target candidate $t = c_i$, e.g., $p(q, c_i) = \sum_j \chi(q, c_i, c_j)$. There has been a large body of work preceding Transformer models, e.g., (Laskar et al., 2020; Tayyar Madabushi et al., 2018; Rao et al., 2016). However, these methods are largely outperformed by the pointwise TANDA model.

Listwise reranking: This approach, e.g., (Bian et al., 2017; Cao et al., 2007; Ai et al., 2018), aims

at learning $p(q, \pi)$, $\pi \in \Pi(\mathcal{A})$, using the information on the entire set of candidates. The loss function for training such networks is constituted by the contribution of all elements of its ranked items.

2.2 Joint Modeling in QA

MR is a popular QA task that identifies an answer string in a paragraph or a text of limited size for a question. Its application to retrieval scenario has also been studied (Chen et al., 2017; Hu et al., 2019; Kratzwald and Feuerriegel, 2018). Their joint modeling aspect regards combining sentences from the same paragraphs. However, the large volume of retrieved content makes their use not practical yet.

Jin et al. (2020) use the relation between candidates in Multi-task learning approach for AS2. They do not exploit transformer models, thus their results are rather below the state of the art. Bonadiman and Moschitti (2020) designed several joint models, which improved early neural networks based on CNN and LSTM for AS2, but failed to improve the state of the art using pre-trained Transformer models.

In contrast to the work above, our modeling is driven by an answer support strategy, where the pieces of information are taken from different documents. This makes our model even more unique; it allows us to design innovative joint models, which are still not designed in any MR systems.

2.3 Fact Verification for Question Answering

Fact verification has become a social need given the massive amount of information generated daily. The problem is, therefore, becoming increasingly important in NLP context (Mihaylova et al., 2018). In QA, answer verification is directly relevant due to its nature of content delivery (Mihaylova et al., 2019). This was explored in MR setting by Wang et al. (2018). Recently, Zhang et al. (2020b) proposed to use potential candidates to extract evidence in ranking.

Zhang et al. (2020a) proposed to fact check for product questions using additional associated evidence sentences. The latter are retrieved based on similarity scores computed with both TF-IDF and sentence-embeddings from pre-trained BERT models. While the process is technically sound, the retrieval of evidence is an expensive process, which is prohibitive to scale in production. We instead address this problem by leveraging the top answer candidates.

3 Baseline Models for AS2

In this section, we describe our baseline models, which are constituted by pointwise, pairwise, and listwise strategies.

3.1 Pointwise Models

One simple and effective method to build an answer selector is to use a pre-trained Transformer model, adding a simple classification layer to it, and fine-tuning the model on the AS2 task. Specifically, $q = \text{Tok}_1^q, \dots, \text{Tok}_N^q$ and $c = \text{Tok}_1^c, \dots, \text{Tok}_M^c$ are encoded in the input of the Transformer by delimiting them using three tags: [CLS], [SEP] and [EOS], inserted at the beginning, as separator, and at the end, respectively. This input is encoded as three embeddings based on tokens, segments and their positions, which are fed as input to several layers (up to 24). Each of them contains sublayers for multi-head attention, normalization and feed forward processing. The result of this transformation is an embedding, \mathbf{E} , representing (q, c) , which models the dependencies between words and segments of the two sentences.

For the downstream task, \mathbf{E} is fed (after applying a non-linearity function) to a fully connected layer having weights: W and B . The output layer can be used to implement the task function. For example, a softmax can be used to model the probability of the question/candidate pair classification, as $p(q, c) = \text{softmax}(W \times \tanh(E(q, c)) + B)$.

We can train this model with log cross-entropy loss: $\mathcal{L} = - \sum_{l \in \{0,1\}} y_l \times \log(\hat{y}_l)$ on pairs of texts, where y_l is the correct and incorrect answer label, $\hat{y}_1 = p(q, c)$, and $\hat{y}_0 = 1 - p(q, c)$. Training the Transformer from scratch requires a large amount of labeled data, but it can be pre-trained using a masked language model, and the next sentence prediction tasks, for which labels can be automatically generated. Several methods for pre-training Transformer-based language models have been proposed, e.g., BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2019), ALBERT (Lan et al., 2020).

3.2 Our joint model baselines

To better show the potential of our approach and the complexity of the task, we designed three joint model baselines based on: (i) a multiclassifier approach (a listwise method), and (ii) a pairwise joint model operating over $k + 1$ candidates, and our adaptation of KGAT model (a pairwise method).

Joint Model Multi-classifier The first baseline is also a Transformer-based architecture: we concatenate the question with the top $k + 1$ answer candidates, i.e., $(q[SEP]c_1[SEP]c_2 \dots [SEP]c_{k+1})$, and provide this input to the same Transformer model used for pointwise reranking. We use the final hidden vector E corresponding to the first input token $[CLS]$ generated by the Transformer, and a classification layer with weights $W \in \mathbb{R}^{(k+1) \times |E|}$, and train the model using a standard cross-entropy classification loss: $y \times \log(\text{softmax}(EW^T))$, where y is a one-hot vector representing labels for the $k + 1$ candidates, i.e., $|y| = k + 1$. We use transformer models such as RoBERTa-Base (or Large), and fine-tuned them with the TANDA approach, i.e., using ASNQ (Garg et al., 2020) as first fine-tuning. The scores for the candidate answers are calculated as $(p(c_1), \dots, p(c_{k+1})) = \text{softmax}(EW^T)$. Then, we rerank c_i according their probability.

Joint Model Pairwise To obtain our second baseline, we concatenate the question with each c_i to constitute the (q, c_i) pairs, which are input to the Transformer, and we use the first input token $[CLS]$ as the representation of each (q, c_i) pair. Then, we concatenate the embedding of the pair containing the target candidate, (q, t) with the embedding of all the other candidates’ $[CLS]$. (q, t) is always in the first position. We train the model using a standard classification loss. At classification time, we select one target candidate at a time, and set it in the first position, followed by all the others. We classify all $k + 1$ candidates and use their score for reranking them. It should be noted that to qualify for a pairwise approach, Joint Model Pairwise should use a ranking loss. However, we always use standard cross-entropy loss as it is more efficient and the difference in performance is negligible.

Joint Model with KGAT Liu et al. (2020) presented an interesting model, Kernel Graph Attention Network (KGAT), for fact verification: given a claimed fact f , and a set of evidences $Ev = \{ev_1, ev_2, \dots, ev_m\}$, their model carries out joint reasoning over Ev , e.g., aggregating information to estimate the probability of f to be true or false, $p(y|f, Ev)$, where $y \in \{\text{true}, \text{false}\}$.

The approach is based on a fully connected graph, G , whose nodes are the $n_i = (f, ev_i)$ pairs, and $p(y|f, Ev) = p(y|f, ev_i, Ev)p(ev_i|f, Ev)$, where $p(y|f, ev_i, Ev) = p(y|n_i, G)$ is the label probability in each node i conditioned on the whole

graph, and $p(ev_i|f, Ev) = p(n_i|G)$ is the probability of selecting the most informative evidence. KGAT uses an edge kernel to perform a hierarchical attention mechanism, which propagates information between nodes and aggregate evidences.

We built a KGAT model for AS2 as follows: we replace (i) ev_i with the set of candidate answers c_i , and (ii) the claim f with the question and a target answer pair, (q, t) . KGAT constructs the evidence graph G by using each claim-evidence pair as a node, which, in our case, is $((q, t), c_i)$, and connects all node pairs with edges, making it a fully-connected evidence graph. This way, sentence and token attention operate over the triplets, (q, t, c_i) , establishing semantic links, which can help to support or undermine the correctness of t .

The original KGAT aggregates all the pieces of information we built, based on their relevance, to determine the probability of t . As we use AS2 data, the probability will be about the correctness of t . More in detail, we initialize the node representation using the contextual embeddings obtained with two TANDA-RoBERTa-base models²: the first produces the embedding of (q, t) , while the second outputs the embedding of (q, c_i) . Then, we apply a max-pooling operation on these two to get the final node representation. The rest of the architecture is identical to the original KGAT. Finally, at test time, we select one c_i at a time, as the target t , and compute its probability, which ranks c_i .

4 Joint Answer Support Models for AS2

We propose the Answer Support Reranker (ASR), which uses an answer pair classifier to provide evidence to a target answer t . Given a question q , and a subset of its top- $k+1$ ranked answer candidates, \mathcal{A} (reranked by an AS2 model), we build a function, $\sigma : Q \times C \times C^k \rightarrow \mathbb{R}$ such that $\sigma(q, t, \mathcal{A} \setminus \{t\})$ provides the probability of t to be correct, where C is the set of sentence-candidates. We also design a multi-classifier MASR, which combines k ASR models, one for each different target answer.

4.1 Answer Support-based Reranker (ASR)

We developed ASR architecture described in Figure 1c. This consists of three main components:

1. a *Pointwise Reranker* (PR), which provides the embedding of the input (q, t) , described in Figure 1a. This is essentially the state-of-the-art

²https://github.com/alexa/wqa_tanda

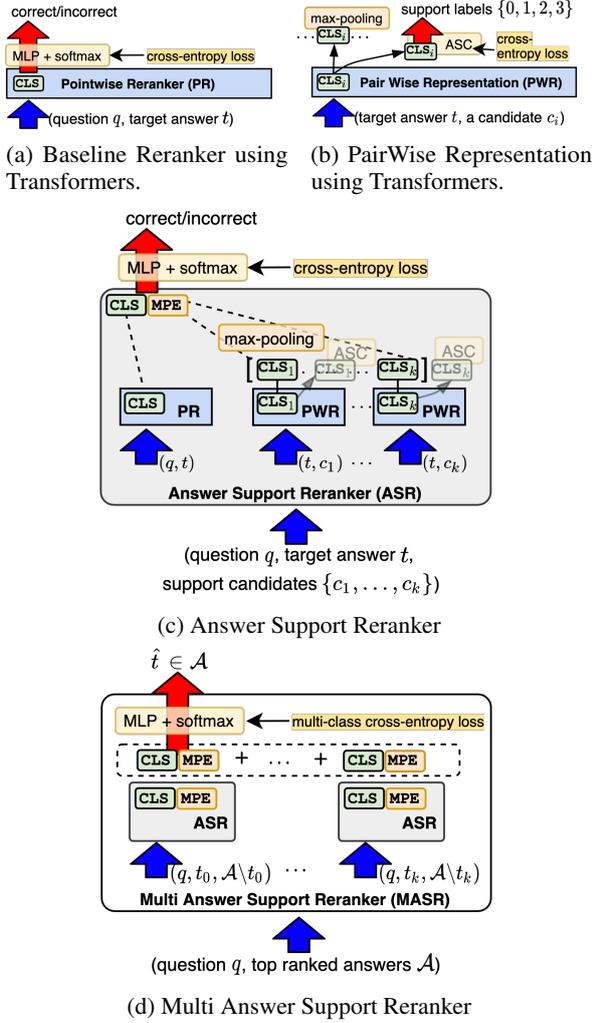


Figure 1: Multi-Answer Support Reranker and its building blocks.

AS2 model based on the TANDA approach applied to RoBERTa pre-trained transformer.

- To reduce the noise that may be introduced by irrelevant c_i , we use the *Answer Support Classifier* (ASC), which classifies each (t, c_i) in one of the following four classes:

- 0 : t and c_i are both correct,
- 1 : t is correct while c_i is not,
- 2 : vice versa, and
- 3 : both incorrect.

This multi-classifier, described in Figure 1b, is built on top a RoBERTa Transformer, which produced a PairWise Representation (PWR). ASC is trained end-to-end with the rest of the network in a multi-task learning fashion, using its specific cross-entropy loss, computed with the labels above.

- The ASR (see Figure 1c) uses the joint representation of (q, t) with (t, c_i) , $i = 1, \dots, k$, where t and c_i are the top-candidates reranked by PR. The k representations are summarized by applying a max-pooling operation, which will aggregate all the supporting or not supporting properties of the candidates with respect to the target answer. The concatenation of the PR embedding with the max-pooling embedding is given as input to the final classification layer, which scores t with respect to q , also using the information from the other candidates. For training and testing, we select a t from the $k + 1$ candidates of q at a time, and compute its score. This way, we can rerank all the $k + 1$ candidates with their scores.

Implementation details: ASR is a PR that also exploits the relation between t and $\mathcal{A} \setminus \{t\}$. We use RoBERTa to generate the $[CLS] \in \mathbb{R}^d$ embedding of $(q, t) = E_t$. We denote with \hat{E}_j the $[CLS]$ output by another RoBERTa Transformer applied to answer pairs, i.e., (t, c_j) . Then, we concatenate E_t to the max-pooling tensor from $\hat{E}_1, \dots, \hat{E}_k$:

$$V = [E_t : \text{Maxpool}([\hat{E}_1, \dots, \hat{E}_k])], \quad (1)$$

where $V \in \mathbb{R}^{2d}$ is the final representation of the target answer t . Then, we use a standard feed-forward network to implement a binary classification layer: $p(y_i | q, t, C_k) = \text{softmax}(VW^T + B)$, where $W \in \mathbb{R}^{2 \times 2d}$ and B are parameters to transform the representation of the target answer t from dimension $2d$ to dimension 2, which represents correct or incorrect labels.

ASC labels There can be different interpretations when attempting to define labels for answer pairs. An alternative to the definition illustrated above is to use the following FEVER compatible encoding:

- 0 : t is correct, while c_i can be any value. This models t correctness, as also an incorrect c_i may provide important context (corresponding to FEVER Support label);
- 1 : t is incorrect and c_i correct. This models t incorrectness, since c_i can provide evidence that t is not similar to a correct answer (corresponding to FEVER Refutation label); and
- 2 : both are incorrect. In this case, nothing can be told (corresponding to FEVER Neutral label).

Dataset	Train			Dev			Test		
	#Q	#A+	#A-	#Q	#A+	#A-	#Q	#A+	#A-
WikiQA	873	1,040	7,632	121	140	990	237	293	2,058
TREC-QA	1,229	6,403	47,014	65	205	912	68	248	1,194
WQA	5,000	42,962	163,289	905	8,179	28,096	1,000	8,256	30,123

Table 2: AS2 dataset statistics

4.2 Multi-Answer Support Reranker (MASR)

ASR still selects answers with a pointwise approach³. This means that we can improve it by building a listwise model that selects the best answer for each question using the information from all target answers. In particular, the architecture of MASR shown in Figure 1d is made up of two parts: (i) a list of ASR containing $k + 1$ ASR blocks, in which each ASR block provides the representation of a target answer t . (ii) A final multiclassifier and a softmax function, which scores each t from $k + 1$ embedding concatenation and selects the one with highest score. For training and testing, we select the t from the $k + 1$ candidates of q based on a softmax output at a time.

Implementation details: The goal of MASR is to measure the relation between $k + 1$ target answers, t_0, \dots, t_k . The representation of each target answer is the embedding $V \in \mathbb{R}^{2d}$ from Equation 1 in ASR. Then, we concatenate the hidden vectors of $k + 1$ target answers to form a matrix $V_{(q,k+1)} \in \mathbb{R}^{(k+1) \times 2d}$. We use this matrix and a classification layer weights $W \in \mathbb{R}^{2d}$, and compute a standard multi-class classification loss:

$$\mathcal{L}_{MASR} = y \times \log(\text{softmax}(V_{(q,k+1)}W^T)), \quad (2)$$

where y is a one-hot-vector, and $|y| = |k + 1|$.

5 Experiments

In these experiments, we compare our models: KGAT, ASR and MASR with pointwise models, which are the state of the art for AS2. We also compare them with our joint model baselines (pairwise and listwise). Finally, we provide an error analysis.

5.1 Datasets

We used two most popular AS2 datasets, and one real world application dataset we built to test the generality of our approach.

WikiQA is a QA dataset (Yang et al., 2015) containing a sample of questions and answer-sentence

³Again, using ranking loss did not provide a significant improvement.

Data split	Supported	Refuted	Not Enough Info
Train	80,035	29,775	35,639
Dev	6,666	6,666	6,666
Test	6,666	6,666	6,666

Table 3: FEVER dataset statistics

candidates from Bing query logs over Wikipedia. The answers are manually labeled. We follow the most used setting: training with all the questions that have at least one correct answer, and validating and testing with all the questions having at least one correct and one incorrect answer.

TREC-QA is another popular QA benchmark by Wang et al. (2007). We use the same splits of the original data, following the common setting of previous work, e.g., (Garg et al., 2020).

WQA The Web-based Question Answering is a dataset built by Alexa AI as part of the effort to improve understanding and benchmarking in QA systems. The creation process includes the following steps: (i) given a set of questions we collected from the web, a search engine is used to retrieve up to 1,000 web pages from an index containing hundreds of millions pages. (ii) From the set of retrieved documents, all candidate sentences are extracted and ranked using AS2 models from (Garg et al., 2020). Finally, (iii) top candidates for each question are manually assessed as correct or incorrect by human judges. This allowed us to obtain a richer variety of answers from multiple sources with a higher average number of answers.

Table 2 reports the corpus statistics of WikiQA, TREC-QA, and WQA⁴.

FEVER is a large-scale public corpus, proposed by Thorne et al. (2018a) for fact verification task, consisting of 185,455 annotated claims from 5,416,537 documents from the Wikipedia dump in June 2017. All claims are labelled as Supported, Refuted or Not Enough Info by annotators. Table 3 shows the statistics of the dataset, which remains the same as in (Thorne et al., 2018b).

5.2 Training and testing details

Metrics The performance of QA systems is typically measured with Accuracy in providing correct answers, i.e., the percentage of correct responses. This is also referred to Precision-at-1 (P@1) in the

⁴The public version of WQA will be released in the short-term future. Please search for a publication with title *WQA: A Dataset for Web-based Question Answering Tasks* on arXiv.org.

RoBERTa Base	WikiQA			TREC-QA			WQA		
	P@1	MAP	MRR	P@1	MAP	MRR	P@1	MAP	MRR
Reranker by Garg et al., 2020	–	0.8890	0.9010	–	0.9140	0.9520	–	–	–
Our Reranker	0.8189 [†]	0.8860	0.8983	0.9118	0.9043	0.9498	0 [†]	0	0
Joint Model Multi-classifier ($k=5$)	0.7819 [†]	0.8542	0.8684	0.8971	0.9052	0.9424	-2.29 [†] %	-1.00%	-1.23%
Joint Model Pairwise ($k=3$)	0.8272 [†]	0.8927	0.9045	0.9559	0.9196	0.9743	2.67 [†] %	0.39%	1.39%
KGAT ($k=2$)	0.8436	0.8991	0.9120	0.9412	0.9155	0.9645	2.10%	0.39%	0.93%
ASR ($k=3$)	†0.8436	0.9014	0.9123	0.9706	0.9257	0.9816	†2.86%	0.86%	1.39%
MASR ($k=3$)	0.8230	0.8891	0.9017	0.9265	0.9200	0.9632	†3.82%	0.70%	1.67%
MASR-F ($k=3$)	0.8272	0.8918	0.9031	0.9412	0.9222	0.9706	2.67%	0.55%	1.47%
MASR-FP ($k=3$)	†0.8436	0.8998	0.9113	0.9559	0.9191	0.9743	†4.96%	0.94%	2.43%

Table 4: Results on WikiQA, TREC-QA and WQA, using RoBERTa base Transformer. The WQA results are reported in terms of difference with Our Reranker baseline (for company policy). [†] is used to indicate that the difference in P@1 between ASR and the other marked systems is statistically significant at 95%.

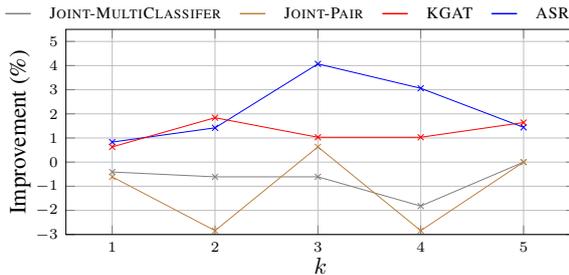


Figure 2: Impact of k on the WQA dev. set

context of reranking, while standard Precision and Recall are not essential in our case as we assume the system does not abstain from providing answers. We also use Mean Average Precision (MAP) and Mean Reciprocal Rank (MRR) evaluated on the test set, using the entire set of candidates for each question (this varies according to the dataset), to have a direct comparison with the state of the art.

Models We use the pre-trained RoBERTa-Base (12 layer) and RoBERTa-Large-MNLI (24 layer) models, which were released as checkpoints for use in downstream tasks⁵.

Reranker training We adopt Adam optimizer (Kingma and Ba, 2014) with a learning rate of $2e-5$ for the transfer step on the ASNQ dataset (Garg et al., 2020), and a learning rate of $1e-6$ for the adapt step on the target dataset. We apply early stopping on the development set of the target corpus for both fine-tuning steps based on the highest MAP score. We set the max number of epochs equal to 3 and 9 for the adapt and transfer steps, respectively. We set the maximum sequence length for RoBERTa to 128 tokens.

KGAT and ASR training Again, we use the Adam optimizer with a learning rate of $2e-6$ for training the ASR model on the target dataset. We

utilize 1 Tesla V100 GPU with 32GB memory and a training batch size of eight. We set the maximum sequence length for RoBERTa Base/Large to 130 tokens and the number of training epochs to 20. The other training configurations are the same of the original KGAT model from (Liu et al., 2020). We use two transformer models for ASR: a RoBERTa Base/Large for PR, and one for ASC. We set the maximum sequence length for RoBERTa to 128 tokens and the number of epochs to 20.

MASR training We use the same configuration of the ASR training, including the optimizer type, learning rate, the number of epochs, GPU type, maximum sequence length, etc. Additionally, we design two different models MASR-F, using an ASC classifier targeting the FEVER labels, and MASR-FP, which initializes ASC with the data from FEVER. This is possible as the labels are compatible.

5.3 Choosing the best k

The selection of the hyper-parameter k , i.e., the number of candidates to consider for supporting a target answer is rather tricky. Indeed, the standard validation set is typically used for tuning PR. This means that the candidates PR moves to the top $k+1$ positions are *optimistically* accurate. Thus, when selecting also the optimal k on the same validation set, there is high risk to overfit the model.

We solved this problem by running a PR version not heavily optimized on the dev. set, i.e., we randomly choose a checkpoint after the standard three epochs of fine-tuning of RoBERTa transformer. Additionally, we tuned k only using the WQA dev. set, which contains $\sim 36,000$ Q/A pairs. WikiQA and TREC-QA dev. sets are too small to be used (121 and 65 questions, respectively). Fig. 2 plots the improvement of four different models, Joint Model

⁵<https://github.com/pytorch/fairseq>

Multi-classifier, Joint Model Pairwise, KGAT, and ASR, when using different k values. Their best results are reached for 5, 3, 2, and 3, respectively. We note that the most reliable curve shape (convex) is the one of ASR and Joint Model Pairwise.

5.4 Comparative Results

Table 4 reports the P@1, MAP and MRR of the rerankers, and different answer supporting models on WikiQA, TREC-QA and WQA datasets. As WQA is an internal dataset, we only report the improvement over PR in the tables. All models use RoBERTa-Base pre-trained checkpoint and start from the same set of k candidates reranked by PR (state-of-the-art model). The table shows that:

- PR replicates the MAP and MRR of the state-of-the-art reranker by Garg et al. (2020) on WikiQA.
- Joint Model Multi-classifier performs lower than PR for all measures and all datasets. This is in line with the findings of Bonadiman and Moschitti (2020), who also did not obtain improvement when jointly used all the candidates altogether in a representation.
- Joint Model Pairwise differs from ASR as it concatenates the embeddings of the (q, c_i) , instead of using max-pooling, and does not use any Answer Support Classifier (ASC). Still, it exploits the idea of aggregating the information of all pairs (q, c_i) with respect to a target answer t , which proves to be effective, as the model improves on PR over all measures and datasets.
- Our KGAT version for AS2 also improves PR over all datasets and almost all measures, confirming that the idea of using candidates as support of the target answer is generally valid. However, it is not superior to Joint Model Pairwise.
- ASR achieves the highest performance among all models (but MASR-FP on WQA), all datasets, and all measures. For example, it outperforms PR by almost 3 absolute percent points in P@1 on WikiQA, and by almost 6 points on TREC from 91.18% to 97.06%, which corresponds to an error reduction of 60%.
- MASR and MASR-F do not achieve better performance than Joint Model Pairwise on WikiQA and TREC, although MASR outperforms all

baselines and even ASR on WQA. This suggests that the significantly higher number of parameters of MASR cannot be trained on small corpora, while WQA has a sufficient number of examples.

- MASR-FP exploiting FEVER for the initialization of ASC performs better than MASR and MASR-F on WikiQA and TREC. Interestingly, it significantly outperforms ASR by 2% on WQA. This confirms the potential of the model when enough training data is available.
- We perform randomization test (Yeh, 2000) to verify if the models significantly differ in terms of prediction outcome. We use 100,000 trials for each calculation. The results confirm the statistically significant difference of ASR and MASR over the baselines, i.e., Our Reranker, Joint Model Multi-classifier, and Joint Model Pairwise, with $p < 0.05$ on both WikiQA and WQA.

5.5 Official State of the art

As the state of the art for AS2 is obtained using RoBERTa Large, we trained KGAT and ASR using this pre-trained language model. Table 5 also reports the comparison with PR, which is the official state of the art. Again, our PR replicates the results of Garg et al. (2020), obtaining slightly lower performance on WikiQA but higher on TREC-QA. KGAT performs lower than PR on both datasets.

ASR establishes the new state of the art on WikiQA with an MAP of 92.80 vs. 92.00. The P@1 also significantly improves by 2%, i.e., achieving 89.71, which is impressively high. Also, on TREC-QA, ASR outperforms all models, being on par with PR regarding P@1. The latter is 97.06, which corresponds to mistaking the answers of only two questions. We manually checked these and found out that these were two annotation errors: ASR achieves perfect accuracy while PR only mistakes one answer. Of course, this just provides evidence that PR based on RoBERTa-Large solves the task of selecting the best answers (i.e., measuring P@1 on this dataset is not meaningful anymore).

5.6 Model Discussion

Table 6 reports the accuracy of ASC inside different models. In ASR, it uses 4-way categories, while in MASR-based models, it uses the three FEVER labels (see Sec. 4.1). ACC is the overall accuracy

RoBERTa Large	WikiQA			TREC-QA		
	P@1	MAP	MRR	P@1	MAP	MRR
Garg et al.,	–	0.9200	0.9330	–	0.9430	0.9740
Our Reranker	0.8724	0.9151	0.9266	0.9706	0.9481	0.9816
KGAT (K=2)	0.8642	0.9094	0.9218	0.9559	0.9407	0.9743
ASR (K=3)	0.8971	0.9280	0.9399	0.9706	0.9488	0.9816

Table 5: Results on WikiQA and TREC-QA, using RoBERTa Large Transformer.

	WikiQA		TREC-QA		WQA	
	ACC	F1	ACC	F1	ACC	F1
ASR	0.59	0.00	0.56	0.80	0.58	0.64
MASR	0.46	0.00	0.45	0.62	0.53	0.61
MASR-F	0.46	0.00	0.64	0.78	0.58	0.68
MASR-FP	0.49	0.37	0.65	0.73	0.59	0.69

Table 6: The Accuracy and F1 of category 0 for ASC

while F1 refers to only the category 0. We note that ASC in MASR-FP achieves the highest accuracy averaged over all datasets. This happens since we pre-fine-tuned it with the FEVER data.

We analyzed examples for which ASR is correct and PR is not. Tab. 7 shows that, given q and $k = 3$ candidates, PR chooses c_1 , a suitable but wrong answer. This probably happens since the answer best matches the syntactic/semantic patterns of the question, which ask for a *type of color*, indeed, the answer offers such type, *primary colors*. PR does not rely on any background information that can support the set of colors in the answer. In contrast, ASR selects c_2 as it can rely on the support of other answers. Its ASC provides an average score for the category 0 (both members are correct) of c_2 , i.e., $\frac{1}{k} \sum_{i \neq 2} \text{ASC}(c_2, c_i) = 0.653$, while for c_1 the average score is significant lower, i.e., 0.522. This provides higher support for c_2 , which is used by ASR to rerank the output of PR.

Tab. 8 shows an interesting case where all the sentences contain the required information, i.e., February. However, PR and ASR both choose answer c_0 , which is correct but not natural, as it provides the requested information indirectly. Also, it contains a lot of ancillary information. In contrast, MASR is able to rerank the best answer, c_1 , in the top position.

6 Conclusion

We have proposed new joint models for AS2. ASR encodes the relation between the target answer and all the other candidates, using an additional Transformer model, and an Answer Support Classifier, while MASR jointly models the ASR representations for all target answers. We extensively tested KGAT, ASR, MASR, and other joint model baselines we designed.

q :	What kind of colors are in the rainbow?
c_1 :	Red, yellow, and blue are called the primary colors.
c_2 :	The order of the colors in the rainbow goes: red, orange, yellow, green, blue, indigo and violet.
c_3 :	The colors in all rainbows are present in the same order: red, orange, yellow, green, blue, indigo, and violet.
c_4 :	A rainbow occurs when white light bends and separates into red, orange, yellow, green blue, indigo and violet.

Table 7: A question with answer candidates ranked by PR; ASR chose c_2 .

q :	What’s the month of Valentine’s day?
c_0 :	Celebrated on February 14 every year, saint Valentine’s day or Valentine’s day is the traditional day on which lovers convey their love to each other by sending Valentine’s cards, sometimes even anonymously.
c_1 :	February is historically chosen to be the month of love and romance and the month to celebrate Valentine’s day.
c_2 :	In order for today to be Valentine’s day, it’s necessary that today is in the month of February.
c_3 :	Every year, Valentine’s day is celebrated on February 14 in many countries around the world.

Table 8: A question with answer candidates $\{c_0, c_1, c_2, c_3\}$ ranked by PR; ASR reranks as $\{c_0, c_3, c_2, c_1\}$; and MASR reranks as $\{c_1, c_3, c_0, c_2\}$; c_1 is the natural correct answer.

The results show that our models can outperform the state of the art. Most interestingly, ASR constantly outperforms all the models (but MASR-FP), on all datasets, through all measures, and for both base and large transformers. For example, ASR achieves the best reported results, i.e., MAP values of 92.80% and 94.88, on WikiQA and TREC-QA, respectively. MASR improves ASR by 2% on WQA, since this contains enough data to train the ASR representations jointly.

References

- Qingyao Ai, Keping Bi, Jiafeng Guo, and W. Bruce Croft. 2018. [Learning a deep listwise context model for ranking refinement](#). *CoRR*, abs/1804.05936.
- Weijie Bian, Si Li, Zhao Yang, Guang Chen, and Zhiqing Lin. 2017. [A compare-aggregate model with dynamic-clip attention for answer selection](#). In *CIKM*, pages 1987–1990. ACM.
- Daniele Bonadiman and Alessandro Moschitti. 2020. [A study on efficiency, accuracy and document structure for answer sentence selection](#). *CoRR*, abs/2003.02349.
- Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. [Learning to rank: from pairwise approach to listwise approach](#). In *Proceedings of the 24th international conference on Machine learning, ICML ’07*, pages 129–136, New York, NY, USA. ACM.

- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading wikipedia to answer open-domain questions](#). *CoRR*, abs/1704.00051.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Siddhant Garg, Thuy Vu, and Alessandro Moschitti. 2020. [TANDA: transfer and adapt pre-trained transformer models for answer sentence selection](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7780–7788. AAAI Press.
- Minghao Hu, Yuxing Peng, Zhen Huang, and Dongsheng Li. 2019. [Retrieve, read, rerank: Towards end-to-end multi-document reading comprehension](#). *CoRR*, abs/1906.04618.
- Zan-Xia Jin, Bo-Wen Zhang, Fang Zhou, Jingyan Qin, and Xu-Cheng Yin. 2020. [Ranking via partial ordering for answer selection](#). *Information Sciences*.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Bernhard Kratzwald and Stefan Feuerriegel. 2018. [Adaptive document retrieval for deep question answering](#). In *EMNLP’18*, pages 576–581.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *ICLR 2020*.
- Md Tahmid Rahman Laskar, Jimmy Xiangji Huang, and Enamul Hoque. 2020. [Contextualized embeddings based transformer encoder for sentence similarity modeling in answer selection task](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 5505–5514, Marseille, France. European Language Resources Association.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020. [Fine-grained fact verification with kernel graph attention network](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7342–7351, Online. Association for Computational Linguistics.
- Tsvetomila Mihaylova, Georgi Karadzhov, Pepa Atanasova, Ramy Baly, Mitra Mohtarami, and Preslav Nakov. 2019. [SemEval-2019 task 8: Fact checking in community question answering forums](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 860–869, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Tsvetomila Mihaylova, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, Mitra Mohtarami, Georgi Karadzhov, and James R. Glass. 2018. [Fact checking in community forums](#). *AAAI-2018*.
- Jinfeng Rao, Hua He, and Jimmy J. Lin. 2016. [Noise-contrastive estimation for answer selection with deep neural networks](#). In *CIKM*, pages 1913–1916. ACM.
- Aliaksei Severyn and Alessandro Moschitti. 2015. [Learning to rank short text pairs with convolutional deep neural networks](#). In *SIGIR’15*.
- Gehui Shen, Yunlun Yang, and Zhi-Hong Deng. 2017. [Inter-weighted alignment network for sentence pair modeling](#). In *EMNLP’17*, pages 1179–1189, Copenhagen, Denmark.
- Harish Tayyar Madabushi, Mark Lee, and John Barn- den. 2018. [Integrating question classification and deep learning for improved answer selection](#). In *COLING’18*, pages 3283–3294.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018a. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018b. [The fact extraction and VERification \(FEVER\) shared task](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- E. Voorhees and D. Tice. 1999. *The TREC-8 Question Answering Track Evaluation*, pages 77–82. Department of Commerce, National Institute of Standards and Technology.
- Mengqiu Wang, Noah A. Smith, and Teruko Mitamura. 2007. [What is the Jeopardy model? a quasi-synchronous grammar for QA](#). In *EMNLP-CoNLL’07*, pages 22–32, Prague, Czech Republic. Association for Computational Linguistics.
- Yizhong Wang, Kai Liu, Jing Liu, Wei He, Yajuan Lyu, Hua Wu, Sujian Li, and Haifeng Wang. 2018. [Multi-passage machine reading comprehension with cross-passage answer verification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1918–1927, Melbourne, Australia. Association for Computational Linguistics.

- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. Wikiqa: A challenge dataset for open-domain question answering. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 2013–2018.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). *CoRR*, abs/1906.08237.
- Alexander S. Yeh. 2000. [More accurate tests for the statistical significance of result differences](#). *CoRR*, cs.CL/0008005.
- Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, and Kyomin Jung. 2019. [A compare-aggregate model with latent clustering for answer selection](#). *CoRR*, abs/1905.12897.
- Wenxuan Zhang, Yang Deng, Jing Ma, and Wai Lam. 2020a. [AnswerFact: Fact checking in product question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2407–2417, Online. Association for Computational Linguistics.
- Yingxue Zhang, Fandong Meng, Peng Li, Ping Jian, and Jie Zhou. 2020b. [Ms-ranker: Accumulating evidence from potentially correct candidates for answer selection](#). *CoRR*, abs/2010.04970.