

CROWN: Conversational Passage Ranking by Reasoning over Word Networks

Magdalena Kaiser, Rishiraj Saha Roy, and Gerhard Weikum

Max Planck Institute for Informatics, Saarland Informatics Campus, Germany
{mkaiser,rishiraj,weikum}@mpi-inf.mpg.de

Abstract. Information needs around a topic often cannot be satisfied in a single turn; users typically ask follow-up questions referring to the same theme. A system must be capable of understanding the conversational context of a request to retrieve correct answers. In this paper, we present our submission to the TREC Conversational Assistance Track (CAST) 2019, in which such a conversational setting is explored. We propose an unsupervised method for conversational passage ranking by formulating the passage score for a query as a combination of similarity and coherence. To be specific, passages are preferred that contain words semantically similar to the words used in the question, and where such words appear close by. We built a word proximity network (WPN) from a large corpus, where words are nodes and there is an edge between two nodes if they co-occur in the same passages in a statistically significant way, within a context window. Our approach, named CROWN, achieved above-average performance on the TREC CAST data with respect to AP@5 and nDCG@1000.

Keywords: Conversations · Passage Ranking · Word Network

1 Introduction

Information needs are usually not one-off: a user who searches for information regarding a specific topic usually asks several questions in a row. Previous turns have an impact on later turns and the system’s answer affects subsequent user queries as well. As a result, questions are often not well-formed and self-contained, but incomplete with ungrammatical phrases and references to previous turns. Thus, a key challenge is to be able to understand context left implicit by the user in their current utterance. However, today’s systems are not capable of answering such questions and there are no resources appropriate for training and evaluating models for such conversational search. The Conversational Assistance Track (CAST)¹ was organized at TREC 2019. The goal was to create a reusable benchmark for open-domain conversational search where answers are passages from large text corpora.

In this work, we describe our submission to TREC CAST 2019. We propose an unsupervised method called CROWN (Conversational Passage Ranking by

¹ <http://www.treccast.ai/>

Reasoning Over Word Networks), in which the passage score for a query is formulated as a combination of similarity and coherence. Similarity between query terms and words in a passage is measured in terms of the cosine similarity of their word embedding vectors. In order to estimate passage coherence, we built a word proximity network (WPN) over a large corpus. At query time, the WPN is used to rank passages preferring those with semantically similar words to the ones appearing in the question and those containing query-relevant term pairs that have an edge in the network. Our CROWN method was able to outperform an Indri baseline on the provided training data and achieved above-average results with respect to AP@5 and nDCG@1000 on the evaluation data. Our code is available on GitHub at <https://github.com/magkai/CROWN>.

2 Related Work

Conversations in search. Conversational questions are often posed to voice assistants. However, current commercial systems cannot handle conversations with incomplete context well. First steps to formalizing conversational search can be found in [15]. Furthermore, a query reformulation approach is described in [18]: an incomplete query is reformulated taking into account the information of previous turns in order to obtain a full-fledged query a standard search engine can deal with. Apart from that, work on search sessions [3, 22, 24] is related to conversational search: information from previous queries from the same session and from click logs is used to better recognize the user’s information need and to improve document ranking. Further works on search sessions focus on query suggestion by using auto-completion logs in addition to click logs [1, 12]. However, in previous works, a ranked list of documents is usually considered as a result for a conversational query, whereas passage-level retrieval, as proposed in TREC CAsT, has not been explored yet.

Conversations in reading comprehension. In machine reading comprehension, answers to questions are text spans in provided paragraphs, like in the SQuAD benchmark [16]. There are several benchmarks available regarding conversational reading comprehension, like QBLink [8], CoQA [17], QuAC [5] and ShARC [19]. A conversational machine reading model is presented, for example, in [25]. Decision rules are extracted from procedural text and reasoning is performed on whether these rules are already entailed by the conversational history or whether the information must be requested from the user. In [14] the pre-trained language model BERT [7] is used to encode a paragraph together with each question and answer in the conversational context and the model predicts an answer based on this paragraph representation. However, these works differ from conversational search, since candidate paragraphs or candidate documents are given upfront.

Conversations over knowledge graphs. Initial attempts in answering conversational questions are also being made in the area of question answering over knowledge graphs (KGs). In [20], the paradigm of sequential question answering over KGs is introduced and a large benchmark, called CSQA, was created for

this task. An unsupervised approach, CONVEX, that uses a graph exploration algorithm is presented in [6] along with another benchmark, named ConvQuestions. Furthermore, approaches based on semantic parsing are presented in [9] and in [21]. These works differ from ours since a knowledge graph is searched for an answer, whereas in TREC CAsT large textual corpora are used as source of answer. Questions over knowledge graphs are mainly objective and factoid, while questions over text corpora have a broader scope. Moreover, answers cannot always be found in KGs due to their incompleteness, whereas the required information can often be located readily in web or news corpora.

3 Task Description

This is the first year of the Conversational Assistance Track in TREC. In this track, conversational search is defined as a retrieval task considering the conversational context. The goal of the task is to satisfy a user’s information need, which is expressed through a sequence of conversational queries (usually ranging from seven to twelve turns in the provided data). Additionally, the topic of the conversation and a description of its content is given. In this year, the conversational topics and turns have been specified in advance. Here is an example for such a conversation taken from the TREC CAsT training data:

Title: Flowering plants for cold climates
Description: You want to buy and take care of flowering plants for cold climates

1. *What flowering plants work for cold climates?*
2. *How much cold can pansies tolerate?*
3. *Does it have different varieties?*
4. *Can it survive frost?*
5. *How do I protect my plants from cold weather?*
6. *How do plants adapt to cold temperature?*
7. *What is the UK hardiness rating for plants?*
8. *How does it compare to the US rating?*
9. *What’s the rating for pansies?*
10. *What about petunias?*

Fig. 1: Sample conversation from the TREC CAsT 2019 data.

As can be seen in the example, subsequent questions contain references to previously mentioned entities and concepts. References like “*it*” in Turn 3, which refer to “*pansy*” or “*What about ...*” in the last turn referring to “*hardiness rating*” cannot be resolved easily. The response from the retrieval system is a ranked list of passages. The passages are short texts (roughly 1-3 sentences each)

Notation	Concept
t, T	Conversational turn t , current turn T
q_t, w_t	Query at turn t (without stopwords), weight for turn t
cq_1, cq_2, cq_3	Conversational query sets
cqw_1, cqw_2, cqw_3	Sets with conversational query weights
$iq_1, iq_2, iq_3, iq_{union}$	Indri query sets
iqw_1, iqw_2, iqw_3	Sets with indri query weights
$G(N, E)$	Word proximity network with nodes N and edges E
NW, EW	Node weights, edge weights
P, P_i, p_{ij}	Set of candidate passages, i^{th} passage, j^{th} token in i^{th} passage
$vec(\cdot)$	Word embedding
$sim(vec(\cdot), vec(\cdot))$	Cosine similarity between word embedding vectors
$NPMI(p_{ij}, p_{ik})$	Normalized point-wise mutual information between two passage tokens
$hasEdge(p_{ij}, p_{ik})$	Returns true if there is an edge between two tokens in the graph
$score_{node}, score_{edge}, score_{indri}$	Similarity score, coherence score, score based on Indri rank
α, β	Threshold for node weights, threshold for edge weights
W	Context window size
h_1, h_2, h_3	Hyperparameters for final score calculation

Table 1: Notation for key concepts in CROWN.

and thus also suitable for voice interfaces or mobile screens. They are retrieved from a combination of three standard TREC collections: MS MARCO Passage Ranking, Wikipedia (TREC CAR) and news (Washington Post).

Thirty example training topics, that have been created manually, are provided by the organizers, as well as relevance judgments on a three-point scale (2: very relevant, 1: relevant, and 0: not relevant) are given for a limited subset. In total, judgments for around 120 questions are specified. The evaluation is performed over 50 different provided topics. Additionally, an Indri baseline using query likelihood is provided. For the baseline run, AllenNLP coreference resolution [11] is performed on the query and stopwords are removed using the Indri stopword list.

4 Method

We now describe CROWN, our unsupervised method for conversational passage ranking. We maximized the passage score for a query that is defined as a combination of similarity and coherence. Intuitively, passages are preferred that contain words semantically similar to the words used in the question and that have such words close to each other. Table 1 gives an overview of notations used for describing our method. Our code as well as some technical details are publicly available at <https://github.com/magkai/CROWN>.

4.1 Building the Word Proximity Network

Word proximity networks have widely been studied in previous literature, for example in [2], where links in a network are defined as significant co-occurrences between words in the same sentence. We chose the MS MARCO Passage Ranking collection as a representative to build the word proximity network for CROWN. We build the graph $G(N, E)$, where nodes N are all words appearing in the

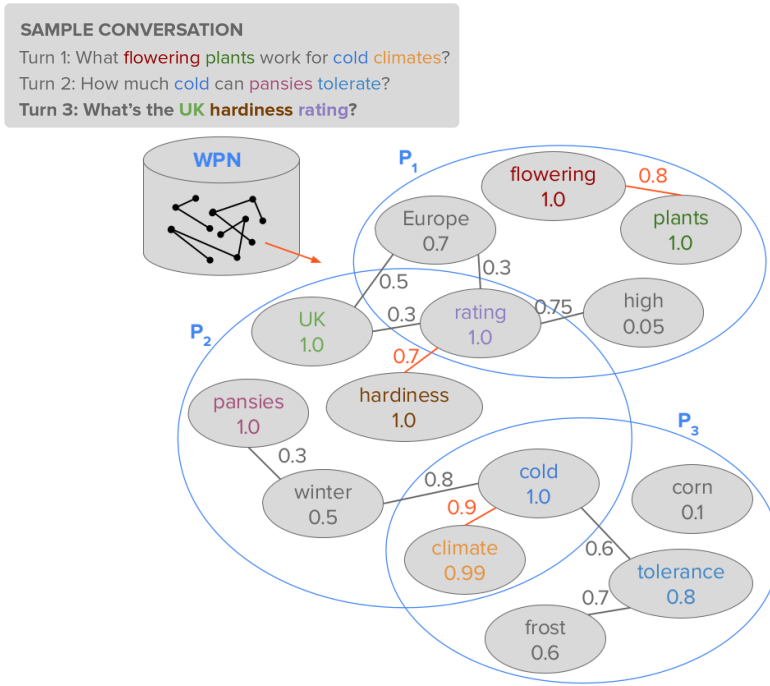


Fig. 2: Sample conversation and word proximity network.

collection (excluding stopwords) and there is an edge $e \in E$ between two nodes if they co-occur in the same passage within a context window W in a statistically significant way. We use NPMI (normalized pointwise mutual information) as a measure of this word association significance, as defined below:

$$npmi(x, y) = \frac{pmi(x, y)}{-\log_2 p(x, y)}$$

where

$$pmi(x, y) = \log \frac{p(x, y)}{p(x) \cdot p(y)}$$

and $p(x, y)$ is the joint probability distribution and $p(x), p(y)$ are the individual distributions over random variables x and y .

While these parts of the network are static and query-agnostic, the network's nodes and edge weights depend on the user input. The NPMI value is used as edge weight between the nodes that are similar to conversational query tokens, whereas node weights are a measure of similarity between conversational query tokens and tokens in the network. In the following sections we will explain the exact weight and score calculations in more detail.

Figure 2 shows a sample conversation, consisting of three turns, together with an excerpt from the word proximity network. In Turn 3 of the conversation “*UK hardiness rating*” is referring to “*pansies*” mentioned in Turn 2. In the sample WPN, nodes consisting of non-stopwords from three candidate passages are displayed. The numeric value in each node is the node weight where the respective color indicates which of the query words is closest to the corresponding word in the passage. For example, the term “*hardiness*” appearing in Passage 2 has a direct match in Turn 3 of the conversation. Therefore, its node weight equals 1. If the similarity is below a threshold, then the corresponding node weight will not be considered further, like “*winter*” or “*Europe*”, marked in grey in the network. The values at the connection of two nodes indicate the edge weights. The edge weights are also only considered if they are above a certain threshold. In our example, only the pairs (“*cold*”, “*climate*”), (“*flowering*”, “*plants*”) and (“*hardiness*”, “*rating*”) appear very frequently together, with an NPMI value of 0.7 or greater, and are therefore considered. They are highlighted in orange in the figure. There are some edge weights which are not considered even though their NPMI values are high. One example for this type is the edge between the pair (“*winter*”, “*cold*”). It is not considered because “*winter*” is not close enough to any query word.

4.2 Formulating the Conversational Query

The three query expansion strategies that worked best in CROWN are described in the following. The conversational query (cq) consists of several queries q_t from different turns t . In the weighted versions, each considered query has a turn weight w_t . When calculating the node weights, the resulting similarity scores are multiplied with the respective turn weight. The conversational query also influences the calculation of the edge weights which we will describe later. The conversational query set cq_1 consists of the current query (q_T) and the first (q_1). No weights are used; therefore, the set of conversational query weights cqw_1 is empty. The second option uses a weighted version (cq_2). It consists of the current, the previous and the first turn, where each turn has a weight in cqw_2 , which is decayed for the previous turn. The last option we consider (cq_3) contains all previous turns, where decaying weights are used for each turn, except for the first and the current one, which receive full weights (cqw_3).

- $cq_1 = \{q_T, q_1\}$, $cqw_1 = \{\}$
- $cq_2 = \{q_T, q_{T-1}, q_1\}$, $cqw_2 = \{w_T, w_{T-1}, w_1\}$, where: $w_1, w_T = 1.0$ and $w_{T-1} = \frac{(T-1)}{T}$
- $cq_3 = \{q_T, q_{T-1}, q_{T-2}, \dots, q_1\}$, $cqw_3 = \{w_T, w_{T-1}, w_{T-2}, \dots, w_1\}$, where $t \in [1, \dots, T]$ **if** ($t == 1 \vee t == T$) **{** $w_t = 1$ **}** **else** **{** $w_t = \frac{t}{T}$ **}**

4.3 Retrieving Candidate Passages

We used the Indri search engine [23] in CROWN to obtain a set of candidate passages P . Our Indri query (iq) also consists of a combination of queries from

different turns. Furthermore, Indri supports weighting of query terms. Here is an example how the weighting of certain words can be done in an Indri query:

```
#weight( 1.0 #combine (survive frost) 0.8 #combine (pansy types) )
```

We were able to produce the best results with the following expansions: In iq_1 , the Indri query consists of the current, the previous and the first turn and no weights are used; iq_2 consists of the current turn, turn $T-1$, turn $T-2$ and the first turn, again without using weights. The weighted version iq_3 uses all previous turns and the corresponding decayed weights can be seen in iqw_3 . Finally, iq_{union} means that three different queries (built from iq_1 , iq_2 and iq_3) are issued to Indri and the union of the resulting passages is used for re-ranking.

- $iq_1 = \{q_T, q_{T-1}, q_1\}$, $iqw_1 = \{\}$
- $iq_2 = \{q_T, q_{T-1}, q_{T-2}, q_1\}$, $iqw_2 = \{\}$
- $iq_3 = \{q_T, q_{T-1}, q_{T-2}, \dots, q_1\}$, $iqw_3 = \{w_T, w_{T-1}, w_{T-2}, \dots, w_1\}$, where $t \in [1, \dots, T]$ in w_t ; **if** $(t == 1 \vee t == T)$ $\{w_t = 1\}$ **else** $\{w_t = \frac{t}{T}\}$
- $iq_{union} = \{iq_1 \cup iq_2 \cup iq_3\}$

4.4 Scoring Candidate Passages

In CROWN, the final score of a passage P_i consists of several components that will be described in the following.

Estimating similarity. The similarity score that is built upon the node weights is calculated in the following way:

$$score_{node}(P_i) = \sum_{j=1}^n \frac{NW(p_{ij})}{\sum_{j=1}^n \mathbb{1}_{C_1}(p_{ij})}$$

where the node weight NW of a token p_{ij} and the condition C_1 will be defined next.

$$NW(p_{ij}) := \mathbb{1}_{C_1}(p_{ij}) \cdot \max_{k \in q_t \in cq} sim(vec(p_{ij}), vec(cq_k)) \cdot w_t$$

where $\mathbb{1}_{C_1}(p_{ij})$ maps to 1 if the condition $C_1(p_{ij})$ is fulfilled, otherwise to 0; $vec(p_{ij})$ is the word embedding vector of the j^{th} token in the i^{th} passage; $vec(cq_k)$ is the word embedding vector of the k^{th} token in the conversational query cq and w_t is the weight of the turn in which cq_k appeared; sim denotes the cosine similarity between the passage token and the query token embeddings. $C_1(p_{ij})$ is defined as

$$C_1(p_{ij}) := \exists cq_k \in cq : sim(vec(p_{ij}), vec(cq_k)) > \alpha$$

which means that condition C_1 is only fulfilled if the similarity between a query word and a word in the passage is above a certain threshold α .

Estimating coherence. Coherence is expressed by term proximity which is reflected in the edge weights. The corresponding score is calculated as follows:

$$score_{edge}(P_i) = \sum_{j=1}^n \sum_{k=j+1}^W \frac{EW(p_{ij}, p_{ik})}{\sum_{j=1}^n \sum_{k=j+1}^W \mathbb{1}_{C_2}(p_{ij}, p_{ik})}$$

The indicator function $\mathbb{1}_{C_2}(p_{ij}, p_{ik})$ maps to 1 if condition $C_2(p_{ij}, p_{ik})$ is fulfilled otherwise to 0. The edge weight EW is defined as:

$$EW(p_{ij}, p_{ik}) := \mathbb{1}_{C_2}(p_{ij}, p_{ik}) \cdot NPMI(p_{ij}, p_{ik})$$

The NPMI value between the tokens is calculated from MS MARCO Passage Ranking as a representative corpus. Condition $C_2(p_{ij}, p_{ik})$ is defined as

$$C_2(p_{ij}, p_{ik}) := C_{21}(p_{ij}, p_{ik}) \wedge C_{22}(p_{ij}, p_{ik})$$

where

$$\begin{aligned} C_{21}(p_{ij}, p_{ik}) &:= hasEdge(p_{ij}, p_{ik}) \wedge NPMI(p_{ij}, p_{ik}) > \beta \\ C_{22}(p_{ij}, p_{ik}) &:= \exists cq_r, cq_s \in cq : \\ &\quad sim(vec(p_{ij}), vec(cq_r)) > \alpha \\ &\quad \wedge sim(vec(p_{ik}), vec(cq_s)) > \alpha \\ &\quad \wedge cq_r \neq cq_s \\ &\quad \wedge \nexists cq_{r'}, cq_{s'} \in cq : \\ &\quad sim(vec(p_{ij}), vec(cq_{r'})) > sim(vec(p_{ij}), vec(cq_r)) \\ &\quad \vee sim(vec(p_{ik}), vec(cq_{s'})) > sim(vec(p_{ik}), vec(cq_s)) \end{aligned}$$

Condition C_{21} assures that there is an edge between the two tokens in the graph and that the edge weight is above a certain threshold β . The second condition, C_{22} , states that there are two non-identical words in the conversational query where one of them is the one that is most similar to p_{ij} (more than any other query token and with similarity above threshold α) and the other word is most similar to p_{ik} .

Estimating priors. We also consider the original ranking received from Indri. In CROWN, this score is defined as:

$$score_{indri}(P_i) = \frac{1}{rank(P_i)}$$

where $rank$ is the rank the passage P_i received from Indri.

Putting it together. The final score for a passage P_i consists of a weighted sum of these three individual scores. More formally:

$$score(P_i) = h_1 \cdot score_{indri}(P_i) + h_2 \cdot score_{node}(P_i) + h_3 \cdot score_{edge}(P_i)$$

where h_1 , h_2 and h_3 are hyperparameters that are tuned using the provided training data.

5 Experimental Setup

5.1 Baseline and Metrics

The provided Indri retrieval model mentioned in Section 3 has been used as the baseline in our experiments. Since responses are assessed using graded relevance, we used nDCG [10] (normalized discounted cumulative gain) and ERR [4] (expected reciprocal rank) as metrics. Furthermore, AP (average precision) is reported on the evaluation data.

5.2 Configuration

Dataset. As mentioned in Section 3, the underlying document collection consists of a combination of three standard TREC collections: MS MARCO, TREC CAR and Washington Post.

Initialization. We used *word2vec* embeddings [13] pre-trained on the Google News dataset and obtained via the python library *gensim*². Furthermore, the python library *spacy*³ has been used for tokenization and stopwords removal. As already mentioned, Indri has been used for candidate passage retrieval. We set the number of retrieved passages from Indri to 1000, so as not to lose any relevant documents. For graph processing, we used the *NetworkX*⁴ python library. The window size W for which word co-occurrences are taken into account is set to three in our graph.

5.3 Submitted Runs

We submitted four runs for the TREC CAsT track. These are described below.

Run 1: mpi-d5_igraph (indri + graph). For our first run, we used the unweighted conversational query cq_1 and the first unweighted option iq_1 for the Indri query. These options performed best in our experiments. For definitions of cq and iq , refer to Section 4.2 and Section 4.3 respectively. The node threshold α is set to 0.7, which means that nodes require having a word embedding similarity to a query token that is greater than 0.7 in order to influence the score calculation. The edge threshold β is set to 0.0 to exclude negative NPMI values. The three hyperparameters are chosen as follows: $h_1 = 0.6$ (indri score), $h_2 = 0.3$ (node score), $h_3 = 0.1$ (edge score).

² gensim: <https://radimrehurek.com/gensim/>

³ spacy: <https://spacy.io/>

⁴ NetworkX: <https://networkx.github.io/>

	mpi-d5_igraph	mpi-d5_intu	mpi-d5_union	mpi-d5_cqw	indri_baseline
nDCG@1000	0.322	0.341	0.195	0.317	0.293
ERR@1000	0.147	0.151	0.038	0.145	0.157

Table 2: Results on training data.




Fig. 3: Turn-wise results for our four runs on evaluation data.

Run 2: mpi-d5_intu (indri-tuned). In our second run, we vary the set of hyperparameters, while the rest stays the same as in Run 1: $h_1 = 0.9$ (indri score), $h_2 = 0.1$ (node score), $h_3 = 0.0$ (edge score). This run gives most emphasis towards the indri score, while coherence in our graph is not considered by giving no weight to the edge score.

Run 3: mpi-d5_union (union of indri queries). Here we use iq_{union} which means that we issue three separate queries to Indri and take the union of all returned passages. However, this leads to three separate Indri rankings which are incomparable. Therefore, we do not consider the indri score in our final score calculation by setting h_1 to 0. Setting $h_2 = 0.6$ (node score) and $h_3 = 0.4$ (edge score) worked best on the training data in this setting. The conversational query and the node threshold are the same as for the previous runs.

Run 4: mpi-d5_cqw (weighted conversational query). In our final run, the conversational query is varied as follows: option cq_2 is used and the node threshold is a bit more restrictive with $\alpha = 0.85$. Apart from that, the parameters are set to the same values as in Run 1.

6 Results and Insights

We present the results of our four runs on the training and the evaluation data.

	mpi-d5_igraph	mpi-d5_intu	mpi-d5_union	mpi-d5_cqw	median	best
Turn 1	0.033	0.033	0.023	0.033	0.035	0.100
Turn 2	0.040	0.042	0.009	0.040	0.031	0.147
Turn 3	0.039	0.040	0.014	0.038	0.035	0.128
Turn 4	0.044	0.052	0.009	0.046	0.043	0.141
Turn 5	0.048	0.053	0.009	0.051	0.041	0.147
Turn 6	0.050	0.059	0.006	0.051	0.039	0.183
Turn 7	0.042	0.051	0.006	0.043	0.038	0.167
Turn 8	0.028	0.032	0.003	0.025	0.029	0.205
All	0.039	0.043	0.010	0.039	0.034	0.145

Table 3: Turn-wise results on evaluation data for AP@5. Above-median values for our submissions are in bold.

	mpi-d5_igraph	mpi-d5_intu	mpi-d5_union	mpi-d5_cqw	median	best
Turn 1	0.497	0.518	0.444	0.497	0.472	0.761
Turn 2	0.448	0.480	0.330	0.446	0.367	0.759
Turn 3	0.486	0.504	0.399	0.479	0.417	0.779
Turn 4	0.438	0.456	0.350	0.436	0.382	0.778
Turn 5	0.425	0.453	0.353	0.410	0.374	0.777
Turn 6	0.454	0.494	0.329	0.458	0.364	0.821
Turn 7	0.463	0.499	0.352	0.456	0.404	0.841
Turn 8	0.374	0.420	0.296	0.376	0.309	0.810
All	0.441	0.470	0.352	0.437	0.362	0.754

Table 4: Turn-wise results on evaluation data for nDCG@1000. Above-median values for our submissions are in bold.

Training data. We compared our runs to the Indri baseline provided by the organizers (see Table 2). Note that for calculating the nDCG and ERR metrics for the training data, only the limited relevance judgements from the manually created conversations have been used. Three of our runs were able to outperform the Indri baseline with respect to nDCG@1000.

Evaluation data. In Tables 3 and 4, the results on the evaluation data for the metrics AP@5 and nDCG@1000 are reported. Average values for each turn (up to Turn 8) and over all turns are displayed. The results for our four runs are reported as well as the median and best turn-wise results over all submissions to the track. Additionally, in Figure 3 the results of our four runs are visualized over eight turns for AP@5 and nDCG@1000.

There seems to be the tendency that the results increase for later turns (up to Turn 6 for AP@5) or do not vary much (up to Turn 7 for nDCG@1000). This means that our method is robust with respect to turn depth and that later turns successfully exploit the information available from previous turns. Three

Turn	Query	Passage Snippet
4	“What makes it so addictive?” (“it”: “ <i>smoking</i> ”, Turn 1)	“Nicotine, the primary psychoactive chemical in cigarettes, is highly addictive.”
2	“What makes it so unusual?” (“it”: “ <i>Uranus</i> ”, Turn 1)	“One fact that is not only unusual, but also makes Uranus markedly different from earth is the angle of its spin axis.”
3	“How much do Holsteins produce?” (“Holsteins”: “ <i>cattle</i> ”, Turn 1) (“produce”: “ <i>milk</i> ”, Turn 2)	“The Holstein-Friesian is the breed of dairy cow most common in [...], around 22 litres per day is average.”
5	“What about in the US?” (“about”: “ <i>physician assistant</i> ”, Turn 1; “ <i>starting salary</i> ”, Turn 4)	“Physician assistant’s starting salary varies from city to city. For instance, in New York [it] is around \$69,559, [...]”
9	“Do NPs or PAs make more?” (“NPs”: “ <i>nurse practitioner</i> ”, Turn 8) (“PAs”: “ <i>physician assistant</i> ”, Turn 1)	“The average salary among all NPs [...] is \$94,881.22 and the average salary among all PAs [...] is \$100,497.78.”

Table 5: Examples for correct answer snippets (rank 1 - 5 in CROWN) for queries from different turns taken from training conversations.

of our runs, namely *mpi-d5_igraph*, *mpi-d5_intu* and *mpi-d5_cqw* achieve above-average performance with respect to both metrics. Our *mpi-d5_union* run does not achieve competitive results probably because the candidate passages which are taken from the union of the three separate Indri retrievals create a pool which is too large for effective re-ranking.

Furthermore, Table 5 shows some exemplary queries taken from the training data, that appear at different turns in the respective conversation, together with passage snippets taken from top-ranked passages by CROWN (rank 1-5). Information from previous turns is required to be able to correctly answer the questions. For example, the query “*What about in the US?*”, asked at Turn 5, needs the additional information “*physician assistants*” and “*starting salary*”, given at Turn 1 and Turn 4 respectively. These are directly matched in the correct answer, resulting in a high node score, and additionally appear next to each other (high edge score).

7 Conclusion

In this work, we presented our unsupervised method CROWN. We showed that using a combination of similarity and coherence for scoring relevant passages is a simple estimate but works quite well in practice. A context window of size three seems to successfully capture significant word co-occurrences. In general, it seemed that giving greater influence to the Indri ranking and giving a higher preference to node weights than edge weights improves results. Regarding query expansion strategies we observed that including only the previous and the first

turns proved to be most beneficial. Weighted turns did not improve the results significantly.

In the future, we would also consider the positions of query terms in passages following the intuition that passages are more relevant in which the query terms appear earlier. Furthermore, we could use contextualized embeddings, like a pre-trained BERT model [7], to deal with polysemy. Yet another possibility would be to use BERT for re-ranking by introducing its ranking results as an additional score.

References

1. Bar-Yossef, Z., Kraus, N.: Context-sensitive query auto-completion. In: WWW (2011)
2. Cancho, R.F.I., Solé, R.V.: The small world of human language. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **268**(1482) (2001)
3. Cao, H., Jiang, D., Pei, J., Chen, E., Li, H.: Towards context-aware search by learning a very large variable length hidden markov model from search logs. In: WWW (2009)
4. Chapelle, O., Metzler, D., Zhang, Y., Grinspan, P.: Expected reciprocal rank for graded relevance. In: *Proceedings of the 18th ACM conference on Information and knowledge management* (2009)
5. Choi, E., He, H., Iyyer, M., Yatskar, M., Yih, W.t., Choi, Y., Liang, P., Zettlemoyer, L.: QuAC: Question answering in context. In: EMNLP (2018)
6. Christmann, P., Saha Roy, R., Abujabal, A., Singh, J., Weikum, G.: Look before you hop: Conversational question answering over knowledge graphs using judicious context expansion. In: CIKM (2019)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL-HLT (2019)
8. Elgohary, A., Zhao, C., Boyd-Graber, J.: A dataset and baselines for sequential open-domain question answering. In: EMNLP (2018)
9. Guo, D., Tang, D., Duan, N., Zhou, M., Yin, J.: Dialog-to-action: conversational question answering over a large-scale knowledge base. In: NeurIPS (2018)
10. Järvelin, K., Kekäläinen, J.: Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)* **20**(4) (2002)
11. Lee, K., He, L., Lewis, M., Zettlemoyer, L.: End-to-end neural coreference resolution. In: EMNLP (2017)
12. Li, L., Deng, H., Dong, A., Chang, Y., Baeza-Yates, R., Zha, H.: Exploring query auto-completion and click logs for contextual-aware web search and query suggestion. In: WWW (2017)
13. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems* (2013)
14. Ohsugi, Y., Saito, I., Nishida, K., Asano, H., Tomita, J.: A simple but effective method to incorporate multi-turn context with BERT for conversational machine comprehension. In: *Proceedings of the First Workshop on NLP for Conversational AI* (2019)
15. Radlinski, F., Craswell, N.: A theoretical framework for conversational search. In: CHIIR (2017)

16. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: Squad: 100,000+ questions for machine comprehension of text. In: EMNLP (2016)
17. Reddy, S., Chen, D., Manning, C.D.: Coqa: A conversational question answering challenge. TACL **7** (2019)
18. Ren, G., Ni, X., Malik, M., Ke, Q.: Conversational query understanding using sequence to sequence modeling. In: WWW (2018)
19. Saeidi, M., Bartolo, M., Lewis, P., Singh, S., Rocktäschel, T., Sheldon, M., Bouchard, G., Riedel, S.: Interpretation of natural language rules in conversational machine reading. In: EMNLP (2018)
20. Saha, A., Pahuja, V., Khapra, M.M., Sankaranarayanan, K., Chandar, S.: Complex sequential question answering: Towards learning to converse over linked question answer pairs with a knowledge graph. In: AAAI (2018)
21. Shen, T., Geng, X., Qin, T., Guo, D., Tang, D., Duan, N., Long, G., Jiang, D.: Multi-task learning for conversational question answering over a large-scale knowledge base. In: EMNLP-IJCNLP (2019)
22. Shen, X., Tan, B., Zhai, C.: Context-sensitive information retrieval using implicit feedback. In: SIGIR (2005)
23. Strohman, T., Metzler, D., Turtle, H., Croft, W.: Indri: A language-model based search engine for complex queries. Information Retrieval - IR (2005)
24. Sun, A., Lou, C.H.: Towards context-aware search with right click. In: SIGIR (2014)
25. Zhong, V., Zettlemoyer, L.: E3: Entailment-driven extracting and editing for conversational machine reading. In: ACL (2019)