

Learning to Search in Long Documents Using Document Structure

Mor Geva

Tel Aviv University

morgeva@mail.tau.ac.il

Jonathan Berant

Tel Aviv University

joberant@cs.tau.ac.il

Abstract

Reading comprehension models are based on recurrent neural networks that sequentially process the document tokens. As interest turns to answering more complex questions over longer documents, sequential reading of large portions of text becomes a substantial bottleneck. Inspired by how humans use document structure, we propose a novel framework for reading comprehension. We represent documents as trees, and model an agent that learns to interleave quick navigation through the document tree with more expensive answer extraction. To encourage exploration of the document tree, we propose a new algorithm, based on Deep Q-Network (DQN), which strategically samples tree nodes at training time. Empirically we find our algorithm improves question answering performance compared to DQN and a strong information-retrieval (IR) baseline, and that ensembling our model with the IR baseline results in further gains in performance.

1 Introduction

Reading comprehension (RC), the task of reading documents and answering questions about their content, has attracted immense attention recently. While early work focused on simple questions and short paragraphs (Hermann et al., 2015; Rajpurkar et al., 2016; Trischler et al., 2017; Onishi et al., 2016), current work is shifting towards more complex questions that require reasoning over long documents (Joshi et al., 2017; Hewlett et al., 2016; Welbl et al., 2017; Kočiský et al., 2017).

Long documents pose a challenge for current RC models, as they are dominated by recurrent neural networks (RNNs) (Chen et al., 2016; Kadlec et al., 2016; Xiong et al., 2017). RNNs process documents token-by-token, and thus using them for long documents is prohibitive. A common solution is to retrieve part of the document with an IR approach (Chen et al., 2017; Clark and Gardner, 2017) or a cheap model (Watanabe et al., 2017), and run an RNN over the retrieved excerpts. However, as documents become longer and questions become complex, two problems emerge: (a) retrieving all the necessary text with a one-shot IR method when performing complex reasoning becomes harder, and thus thousands of tokens are retrieved (Clark and Gardner, 2017). (b) Running even a cheap model over the document in its entirety becomes expensive (Choi et al., 2017).

Humans, in lieu of a mental inverted index, use document structure to guide their search for answers. E.g., the answer to “*What high school did Leonard Cohen go to?*” is likely to appear in “*Early life*”, while the answer to “*How hot is it in Melbourne in July?*” is likely to appear in “*Climate*”. In this work we investigate whether we can train a model to navigate through a document using its structure and find the answer, while reading only a small portion of the entire document.

We represent documents as trees and train an agent that navigates through the document tree until returning a final answer. Figure 1 illustrates this process. Our agent reads the question “*The vacation destinations of Pattaya and Phuket are in which country?*” and starts navigation at the title of the document. After reading a paragraph, and skipping “*History*”, it drills down to “*Geography*” until finally halting at a paragraph that specifies the answer (“*Thailand*”). The agent observes at each step only a glimpse of the local text to determine its next action, which can be movement to a tree node, answering

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

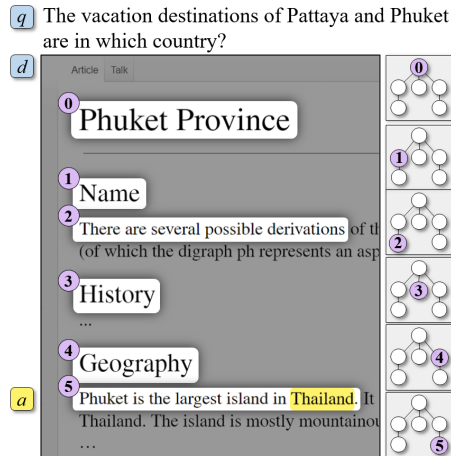


Figure 1: An overview of our framework: an agent answers the question q for a document d by starting at the title, performing navigation actions until reaching the relevant paragraph, extracting the answer a , and then stopping.

the question with a more expensive RC model, or terminating navigation. Thus, the agent consumes only a small fraction of the entire document.

Our training data is question-document-answer triplets, without gold navigation paths, and thus we train our model with the Deep Q-Network (DQN) algorithm. Because the dataset is biased towards answers appearing at the beginning of the document, the algorithm tends to stop early and does not explore the document well. To overcome this challenge, we propose DOCQN: a variant of DQN for tree navigation that improves exploration by sampling nodes from multiple parts of the tree.

We evaluate the ability of our agent to navigate to paragraphs containing the answer on a variant of TRIVIAQA (Joshi et al., 2017) and find that: (a) DOCQN navigates better than DQN in documents both quantitatively and qualitatively. (b) While DOCQN observes only 6% of the document tokens, it outperforms an IR method in end-to-end QA performance. (c) An ensemble of DOCQN and IR substantially improves both navigation and end-to-end QA performance over the ensemble components.

To summarize, in this paper we ask: can an agent use document structure and learn to find answers for complex questions in long documents? We propose a new model and training algorithm that overcomes an inherent bias in the data, answering the aforementioned question in the affirmative. Our code and dataset are available at <https://github.com/mega002/DocQN>.

2 Problem Overview

We work in the traditional RC setup, where we are given question-document-answer triplets $\{(q_i, d_i, a_i)\}_{i=1}^N$ as a training set, and aim to learn a function that finds the answer for an unseen question-document pair. Unlike prior work, we assume documents are trees, where every tree node u corresponds to a structural element and is labeled with text $l(u)$. Specifically, the root is labeled with the document title, sections and subsections are labeled by their title, and paragraphs and sentences are labeled by the text they contain. In addition, we order all non-sentence tree nodes by a pre-order traversal (which corresponds to the linear order of text in the document), and denote the index of a node u by $n(u)$. For sentence nodes, $n(u)$ is the index of their parent (a paragraph).

Figure 1 shows an example tree, where for each node we show the relevant structural element and index (sentence nodes are not shown in the figure).

With this document representation, answering questions can be viewed as a Markov Decision Process (MDP), where in each state the agent is located in a particular tree node, actions allow movement through the document tree, answering the question with a RC model, or stopping, and a reward is based on whether the agent locates a node that contains the answer.

Our agent interleaves actions that *navigate* in the document, with an action that runs an RC model in a certain document position and *extracts an answer*. Thus, the agent can decide to continue navigation after extracting a certain answer. This is strictly more expressive than existing approaches that combine

		TRIVIAQA	TRIVIAQA-NOP
Train	Questions	61,888	57,220
	Documents	110,647	99,315
Development	Questions	7,993	7,336
	Documents	14,229	12,706
Test	Questions	7,701	6,507
	Documents	13,661	10,481

Table 1: Data statistics for TRIVIAQA vs. TRIVIAQA-NOP.

Average number of tokens	5590.7
Average number of tree nodes	332.2
Average number of high-level sections	6.6

Table 2: Data statistics for the evidence documents of TRIVIAQA-NOP.

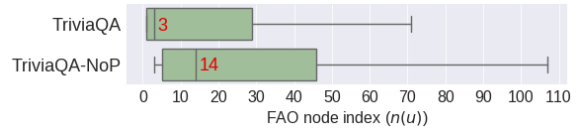


Figure 2: FAO node index distribution of TRIVIAQA and TRIVIAQA-NOP (median values in red).

	TRIVIAQA	TRIVIAQA-NOP
Questions	86.7%	83.8%
Question-document pairs	69.7%	66.9%

Table 3: Portion of answerable samples in a random subset of the training set, which contains 278 questions and 475 question-document pairs.

IR with RC, where some text is retrieved exactly once before applying an RC model. As RC shifts to reasoning over complex questions, navigating and reading multiple parts of the document will become necessary. We show in Section 5 this approach indeed improves QA performance.

3 Data

To test our framework, we capitalize on the recently-released TRIVIAQA dataset, which contains question-answer pairs, along with a small set of documents that (in almost all cases) contain the answer. TRIVIAQA is suitable for our purposes as it is a large scale dataset, where questions are relatively complex and documents are fairly long. The dataset includes only raw text, and thus for every evidence document, we built a tree representation by extracting the html metadata from the corresponding Wikipedia page, and constructing the document structure from it.

Because our goal is to investigate whether a model can learn to search through a document, it is important that a non-negligible fraction of the questions require navigation through the document. However, in Wikipedia each document starts with a preface that summarizes the document, and thus often contains the answer. Consequently, a model that ignores the question and document and always stops in the first paragraph is likely to obtain good performance. Figure 2 shows the distribution of the first answer occurrence (FAO) in a document in TRIVIAQA over the node indices $n(u)$ (x -axis), where all tree nodes, except sentences are considered (see Section 2). We find that in most question-document pairs the FAO is in the first few paragraphs, and that in 60% of the cases it is in the preface section.

To alleviate this heavy bias, we derive a new dataset, termed TRIVIAQA-NOP, where we remove the preface section from every document. After removing the preface, 2,144 out of 77,582 (2.8%) questions and 6,124 out of 138,537 (4.4%) question-document pairs are left without an answer and are removed. To further reduce the number of cases where an answer can not be inferred from a document, we drop question-document pairs where: (a) the answer appears only in titles; (b) the answer is a single-character; (c) the FAO node index is > 700 (in most cases the answer is an item in a list). Finally, the dataset includes 91.6% of the questions and 88.4% of the question-document pairs from the original dataset. We provide full statistics on the dataset in Tables 1 and 2.

To verify that questions remain answerable in TRIVIAQA-NOP after removing the preface, we perform manual analysis on a random sample of 278 questions and 475 documents from the training set (Table 3). We find that the portion of answerable questions and question-document pairs remains high and is reduced by less than 3% in comparison to TRIVIAQA. This demonstrates that indeed, for most questions and documents, the context necessary for answering the question also appears in the document body and not only in the preface.

Figure 2 shows the FAO node index distribution in TRIVIAQA-NOP. We observe, compared to TRIVIAQA, that the first occurrence of an answer is much more spread out across the document, and that the median increases from 3 to 14, which will require more navigation from the agent. However, even in TRIVIAQA-NOP answers tend to appear at the beginning of the document, because document content is usually organized by importance. This bias results in an exploration challenge for our training algorithm,

	Example
o	"Phuket Province Name"
ϕ_n^1 : height	2
ϕ_n^2 : depth	1
ϕ_n^3 : h_dist_start	0
ϕ_n^4 : h_dist_end	2
ϕ_n^5 : parent.h_dist_start	0
ϕ_n^6 : parent.h_dist_end	0
ϕ_n^7 : navigation step	1

Table 4: An example observation o and list of navigation features ϕ_n for node 1 in Figure 1. The features `height` and `depth` correspond to distance from the farthest leaf and root, respectively (`height` is 2 since sentence nodes are omitted from the figure). `h_dist_start` and `h_dist_end` measure the horizontal distance from the first and last child of the node’s parent. `navigation step` is a counter for the number of performed steps.

which we will address in Section 4.

4 Method

In this section, we describe a model for the agent and a training algorithm based on DQN (Mnih et al., 2015). Specifically, we introduce a tree-sampling strategy, which addresses the exploration challenge stemming from the bias towards answers early in the document.

4.1 Framework

We represent the MDP as a tuple $(\mathcal{S}, \mathcal{A}, R, T)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $R(s, a)$ is a reward function, and $T(s, a)$ is a deterministic transition function. Our model implements an action-value function $Q(s, a)$, which takes a state $s \in \mathcal{S}$ and returns a value for every action $a \in \mathcal{A}$. This function defines a policy $\pi(s) = \arg \max_{a \in \mathcal{A}} Q(s, a)$. We now describe the state space, actions and reward function.

States Given a tree node u , a state is a tuple $s^u = (q, o, z, \phi_n, \phi_z)$, where q is the question, o is an observation, z is an answer prediction, and ϕ_n, ϕ_z are navigation and answer prediction features. An observation $o = (o_1, \dots, o_{|o|})$ is a sequence of tokens produced by recursively concatenating the first k tokens of text in the label $l(u)$ to the observation of u ’s parent. An answer prediction $z = (z_1, \dots, z_{|z|})$ is the sequence of tokens that were extracted by an RC model, if an RC model was already run on $l(u)$ (and a null token otherwise). The answer prediction features $\phi_z = (z_e, z_l, z_n)$ provide information on the distribution over answer spans provided by the RC model, which reflects its confidence: z_e is the entropy of the distribution, z_l is the logit value for z , and z_n is the number of tokens in $l(u)$. Navigation features ϕ_n provide information on the relative location of u in the document. An example observation and full list of navigation features are given in Table 4.

Note that the state s does not depend on the history of visited tree nodes.¹ While incorporating history could be beneficial, a memory-less model enables us to explore tree-sampling strategies, which is important for training (Section 4.2).

Actions We define the following set of actions \mathcal{A} . Let u be a node with an ordered list of children (v_1, v_2, v_3) , and w be a child of v_2 . We define five movement actions (Figure 3), where `DOWN` moves from u to its first child v_1 , `RIGHT` moves from v_2 to v_3 , and `LEFT` moves from v_2 to v_1 . Because moving upwards reaches a node we already visited, we define `UPR`, which moves from w to v_3 , and `UPL`, which moves from w to v_1 . If an illegal action is chosen (e.g., `LEFT` from v_1), then the agent stays in its current position. The action `ANSWER` returns an answer (and a distribution over spans) by running a RC model on $l(u)$, unless u is a sentence, in which case it is run on the paragraph containing u . After `ANSWER`, the agent can resume navigation. The action `STOP` also returns the answer given the current node u , but also terminates navigation.

Reward Our goal is to develop an agent that can navigate in the document, and thus we define the reward based on whether the agent stops in a node that contains the gold answer text (this is noisy,

¹except for `navigation step`, which can be approximated by the shortest path from the root to any node.

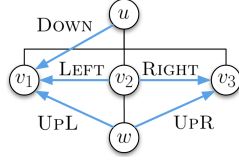


Figure 3: Movement actions in our environment.

because the answer might be there sporadically). While a simple reward would be an indicator for whether the agent stopped at a correct node, such a reward would not capture the proximity of the agent to the answer. Moreover, we would like to consider the overall document length, rewarding successful navigations in long documents. Therefore, we define the following reward:

$$r(s, a) = \begin{cases} 2 & a = \text{STOP}, |n(u) - n(u^*)| = 0 \\ 1 - \frac{|n(u) - n(u^*)|}{\max_u n(u)} & a = \text{STOP}, |n(u) - n(u^*)| > 0 \\ -0.06 & a = \text{ANSWER} \\ -0.02 & a \notin \{\text{ANSWER}, \text{STOP}\} \end{cases}$$

where u is the node where the agent is located and u^* is the closest tree node that contains the answer ($n(u)$ is the node index as defined above). Thus, when stopping, the reward is proportional to the distance to the closest answer location given the document length. An additional reward is given if navigation is successful, and a penalty is given for any other action, to encourage shorter trajectories. We further penalize the ANSWER action to discourage frequent usage of the RC model.

4.2 DOCQN: DQN with Tree Sampling

Training the navigation model is based on DQN (Mnih et al., 2015). In DQN, at every step an agent at state s_t selects an action a_t using ϵ -greedy policy, given the current action-value function $Q_\theta(s_t, a_t)$ parameterized by θ . The agent observes a reward r_t and a state $s_{t+1} = T(s_t, a_t)$ and adds a transition (s_t, a_t, r_t, s_{t+1}) to a replay memory buffer \mathcal{D} that holds a large number of recent transitions. The parameters are then optimized so that the action-value function matches better the observed reward. This is done by sampling a batch of random transitions from \mathcal{D} and minimizing the regression loss

$$(r_t + \gamma \max_{a'} Q_{\hat{\theta}}(s_{t+1}, a') - Q_\theta(s_t, a_t))^2, \quad (1)$$

where $\hat{\theta}$ are the parameters of a *target network*, which is a periodic copy of θ that is not optimized, and γ is a discount factor.

We also add some of the recent enhancements to DQN (Hessel et al., 2017), which have proved to be useful in our setup. Specifically we implement *Double Q-Learning* (van Hasselt et al., 2016), *Prioritized Experienced Replay* (Schaul et al., 2015), and *Dueling Networks* (Wang et al., 2016).

Reducing bias with DOCQN The DQN algorithm contains episodes, where in each episode the agent is placed at an initial state s_0 , from which it starts taking actions. In our setup, this state corresponds to the root of the document tree. Because the data has a bias towards answers appearing at the beginning of the document, the agent learns that stopping early improves the reward and is stuck at a local minimum, where it ceases exploration. Examining Figures 2 and 8, we observe that DQN learns to stop very early in the document compared to the FAO node index distribution of TRIVIAQA-NOP, amplifying the bias in the data.

To address this issue, we suggest DOCQN, a variant of DQN aimed at increasing exploration when navigating in a document structure. DOCQN capitalized on two properties. First, DQN is an off-policy algorithm that trains from transitions (s_t, a_t, r_t, s_{t+1}) . Second, our model is memory-less, and thus we can sample any node u and compute the corresponding state s^u . Therefore, we modify DQN, and instead of initializing every episode with s_0 and performing a sequence of actions, we sample states from distributions that explore the document better. By exploring transitions from across the document, the model learns from more distant parts of the document.

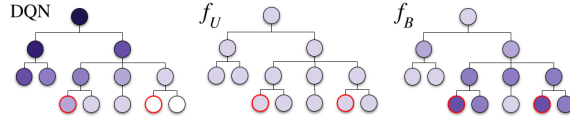


Figure 4: Illustration of the different state distributions for a specific documents tree. Nodes with darker color have higher sampling probability, and nodes marked in red are paragraph nodes containing the answer.

Algorithm 1: **DOCQN**

```

1: Let  $f$  be a distribution over tree nodes
2: Initialize replay memory  $\mathcal{D}$  and parameters  $\theta, \hat{\theta}$ 
3: for episode = 1,  $M$  do
4:   if random() <  $\epsilon_s$  then
5:     for  $i = 1 \dots K$  do
6:       sample node  $u \sim f$  and generate  $s^u$ 
7:       sample action  $a$  from  $s^u$  ( $\epsilon$ -greedily).
8:        $r \leftarrow R(s^u, a), s' \leftarrow T(s^u, a)$ 
9:       store  $(s^u, a, r, s')$  in  $\mathcal{D}$ 
10:   else
11:     Initialize start state  $s_0$ 
12:     for  $t = 0 \dots T - 1$  do
13:       sample action  $a$  from  $s_t$  ( $\epsilon$ -greedily).
14:        $r \leftarrow R(s_t, a), s_{t+1} \leftarrow T(s_t, a)$ 
15:       store  $(s_t, a, r, s_{t+1})$  in  $\mathcal{D}$ 
16:   UPDATEPARAMS( $\theta, \hat{\theta}, \mathcal{D}$ ) (Equation 1)

```

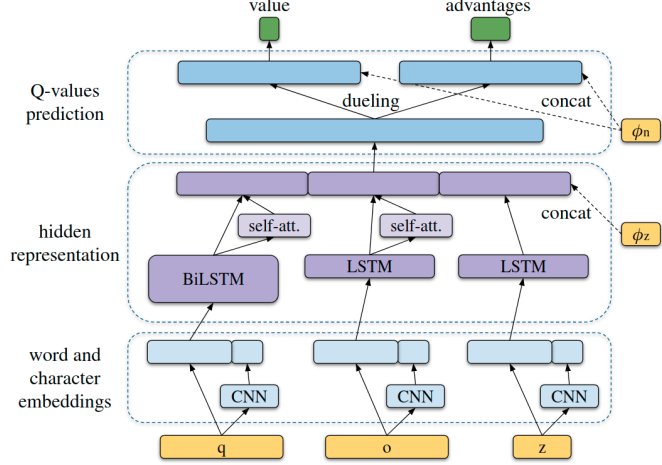


Figure 5: High-level overview of the network architecture.

Algorithm 1 provides the details of DOCQN. The algorithm initializes the replay memory buffer, the model and target network parameters, and chooses a distribution f over tree nodes. At the beginning of every episode, with probability ϵ_s (line 4),² we sample K state transitions using f and store them in \mathcal{D} , and with probability $(1 - \epsilon_s)$ (line 10), we start at the initial state s_0 and sample a trajectory as in DQN. Parameters are updated as in DQN by sampling transitions from the replay buffer and minimizing Equation 1. If f samples nodes from various locations in the document, the algorithm will explore better and will not getting stuck at its beginning. We consider the following instantiations of f :

1. f_U : *Uniform sampling* over nodes, except we discourage sampling sentences by uniformly sampling with probability 0.2 a leaf (sentence), and probability 0.8 an inner node.
2. f_B : *Backward sampling*: We uniformly sample a paragraph node p that contains the gold answer, then uniformly a number of actions $B \in \{1, 2, 3\}$, and perform B random movement actions from p to output a node. This results in a node that is “close” to the answer a , and can be viewed as similar to bi-directional search or backward search (Lao et al., 2015).

Figure 4 shows the distribution over tree nodes for a specific document tree where the answer appears in two paragraph nodes (ignoring sentence nodes). We see that sequential sampling (as in DQN) puts most of the probability mass close to the root, uniform sampling is uniform (across paragraphs, as sentence nodes are omitted), and backward sampling is concentrated close to answer nodes. This illustrates how different distributions result in different exploration strategies.

Network Architecture We briefly describe the neural architecture (Figure 5) and provide full details in the supplementary material. As explained in Section 4.1, the input to the network is the state s , which comprises the question tokens q , observation tokens o , answer prediction tokens z and features ϕ_n, ϕ_z . The question, observation and answer prediction tokens are encoded with pre-trained word embeddings and trained character embeddings, where character embeddings are followed by a convolutional layer with max pooling, yielding a single vector per token. Each token is then represented by concatenating the word embedding with the character embedding.

Question tokens, observation tokens, and answer to are then fed into a BiLSTM and LSTM (Hochreiter and Schmidhuber, 1997) respectively and the LSTM outputs are compressed to a single vector through

² ϵ_s is annealed from 1 to 0.5 during training.

self attention (Cheng et al., 2016), resulting in one vector for q and one for o . Answer prediction tokens are fed into an LSTM, where the last hidden state is concatenated to the features ϕ_z , thus creating a third vector for the answer prediction.

We concatenate these three vectors, and pass them through a one layer feed-forward network that then branches to two networks according to the Dueling DQN architecture (Wang et al., 2016). In each branch we also concatenate the navigation features ϕ_n . One branch predicts the value of the state $V_\theta(s) \in \mathbb{R}$, and the other branch predicts the advantage of every action $A_\theta(s, a) \in \mathbb{R}$ for every possible action a . The output of the network is $Q_\theta(s, a) = V_\theta(s) + (A_\theta(s, a) - \frac{1}{A} \sum_{a'} A_\theta(s, a'))$ as in Wang et al. (2016).

5 Experimental Evaluation

Our experimental evaluation aims to answer the following questions: (a) Can document structure be used to learn to navigate to an answer in a document? (b) How does DOCQN compare to DQN? (c) How does DOCQN compare to IR methods that observe the entire document?

5.1 Experimental Setup

We evaluate on TRIVIAQA-NOP. Because our focus is on the navigation ability of the agent, we train a single RC model and fix it in all experiments. Specifically, we download RASOR (Lee et al., 2017; Salant and Berant, 2018),³ and exactly follow the procedure described by the authors of TRIVIAQA for training a RC model, i.e., we train RASOR on the first 400 tokens of each document in TRIVIAQA-NOP. As a sanity check for our RC model, we also train and evaluate RASOR on the original TRIVIAQA dataset. Indeed, RASOR obtains 48.6% EM and 53.4% F1, which is substantially higher than the baseline reported by Joshi et al. (2017).

Evaluation We use two metrics: First, we measure *navigation accuracy*, i.e., for a question-document pair, whether a method returns text containing a gold answer (if the agent stops at a sentence node we evaluate the encompassing paragraph). Because questions in TRIVIAQA-NOP often have more than one evidence document, we also measure *aggregated navigation accuracy*, where we give credit if the agent navigated correctly in any of the documents. This gives performance assuming an oracle that always chooses the best document for the question. Because the test set in TRIVIAQA is hidden, we evaluated navigation accuracy on the development set only.

In addition, we measure end-to-end QA performance with the official Exact Match (EM) and F1 metrics. The EM metric measures the percentage of predictions that match exactly any of the answer aliases, and the F1 metric measures the average overlap between the prediction and answer. To aggregate evidence from multiple documents we follow Joshi et al. (2017) and define the score of an answer to be the sum of probabilities for that answer across all documents.

Models We compare the following models:

- **DOCQN**: This is our main model, where we use the sampling distribution $f_{U+B} = \frac{1}{2}f_U + \frac{1}{2}f_B$.
- **DQN**: DQN algorithm without state sampling.
- **{DOCQN|DQN}-COUPLED**: A less expressive version of the model, where the actions ANSWER and STOP are coupled, i.e., the agent can use the RC model exactly once and then stops. This is similar to a setup where retrieval is performed once and not interleaved with navigation.
- **RANDOMWALK** An agent that selects an action at each step uniformly at random.
- **RANDOMPARA** An agent that randomly selects a non-sentence tree node.
- **DOC-TF-IDF**: An IR baseline, where a paragraph is selected based on its tf-idf score. We implement the tf-idf scheme of DOCQA (Clark and Gardner, 2017), which is a high-performing system on TRIVIAQA. Here, idf is computed for each paragraph in the context of the current document, and the paragraph with highest cosine similarity to the question is chosen. Note that DOC-TF-IDF processes the entire document (up to tens of thousands of tokens), while DOCQN processes only a small fraction.
- **TF-IDF**: Vanilla tf-idf, where the idf score is computed from all documents in TRIVIAQA-NOP.

³<https://github.com/shimisalant/RaSoR>

- **Ensemble:** We ensemble $\text{DOCQN } f_{U+B}$ with DOC-TF-IDF in two ways: (1) for finding the final answer, we aggregate scores as described above, that is, we sum the probabilities from both models over all documents and choose the answer with highest score. (2) For navigation, we simply tune on the development set a threshold l , where we take the prediction of $\text{DOCQN } f_{U+B}$ if it stopped at a node with index $\leq l$ and use DOC-TF-IDF otherwise.
- **READTOP:** Following Joshi et al. (2017), a model that runs the RC model on the first 800 tokens of the document. Note that running the RC model on the text retrieved by DOCQN involves consuming far fewer tokens, namely, RASOR consumes only 160 tokens on average.

We report the value of all hyper-parameters in the supplementary material (all fixed without tuning).

5.2 Results

Step	Node	Observation Tokens	Answer	Action
0	0	['Wednesday']		DOWN
1	1	['Wednesday', 'Etymology']		DOWN
2	2	['Wednesday', 'Etymology', 'See', 'Names', 'of', 'the', 'days', 'of', 'the', 'week', 'for', 'more', 'on', 'naming', 'conventions', '']]		UPR
3	25	['Wednesday', 'Religious', 'observances']		DOWN
4	26	['Wednesday', 'Religious', 'observances', 'The', 'Creation', 'narrative', 'in', 'the', 'Hebrew', 'Bible', 'places', 'the', 'creation', 'of', 'the', 'Sun', 'and', 'Moon', 'on', 'the', 'fourth', 'day']		ANS
5	26	['Wednesday', 'Religious', 'observances', 'The', 'Creation', 'narrative', 'in', 'the', 'Hebrew', 'Bible', 'places', 'the', 'creation', 'of', 'the', 'Sun', 'and', 'Moon', 'on', 'the', 'fourth', 'day']	Sun and Moon	UPR
6	37	['Wednesday', 'Cultural', 'usage']		DOWN
7	38	['Wednesday', 'Cultural', 'usage', 'In', 'Hindu', 'mythology', 'is', 'the', 'god', 'of', 'Mercury', 'planet', 'is', 'mid-week', 'Wednesday', 'and', 'of', 'Merchants']		ANS
8	38	['Wednesday', 'Cultural', 'usage', 'In', 'Hindu', 'mythology', 'is', 'the', 'god', 'of', 'Mercury', 'planet', 'is', 'mid-week', 'Wednesday', 'and', 'of', 'Merchants']	Budha	RIGHT
9	39	['Wednesday', 'Cultural', 'usage', 'In', 'the', 'folk', 'rhyme', 'Wednesday', 'is', 'child', 'is', 'full', 'of', 'woe', 'reciting', 'the', 'days', 'of', 'the', 'week']		ANS
10	39	['Wednesday', 'Cultural', 'usage', 'In', 'the', 'folk', 'rhyme', 'Wednesday', 'is', 'child', 'is', 'full', 'of', 'woe', 'reciting', 'the', 'days', 'of', 'the', 'week']	full of woe	STOP

Figure 6: Navigation example of DOCQN.

Tables 5 and 6 show the results of our experiments. Focusing on navigation accuracy, we see that randomly walking or choosing a paragraph yields low performance. Vanilla TF-IDF performs considerably better than the random baselines, but is outperformed by all other models. Comparing DQN to DOCQN, we see that DOCQN outperforms DQN. Allowing DQN and DOCQN to run the RC model during navigation improves their performance, but the accuracy gain is substantially higher for DOCQN (2.3% accuracy and 3.1% aggregated accuracy) than for DQN (0.4% accuracy and 0.7% aggregated accuracy). DOC-TF-IDF, which has access to the entire document outperforms DOCQN, which consumes on average 6% of the entire document. Nonetheless, the two models obtain the same aggregated accuracy. This good performance of DOC-TF-IDF shows that in TRIVIAQA-NOP answer paragraphs share a lot of lexical material with the question. Importantly, an ensemble of DOC-TF-IDF with DOCQN substantially improves the overall performance, reaching accuracy of 35% and aggregated accuracy of 48%. This is

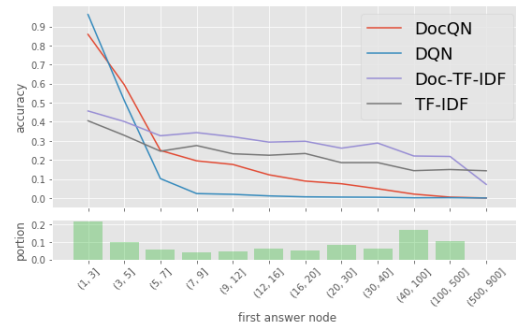


Figure 7: Model performance (top) and portion of samples in the development set (bottom) as functions of the node index of the FAO.

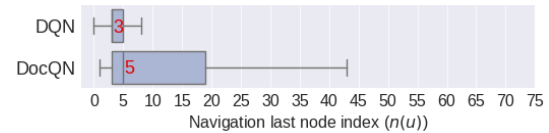


Figure 8: Distribution of node index at navigation stopping point (median value in red) for DQN and DOCQN.

	Dev.	Dev. Agg.
RANDOMWALK	1.8	3.1
RANDOMPARA	13.4	20.9
DQN-COUPLED	26.8	38.8
DQN	27.2	39.5
DOCQN-COUPLED	28.1	40.5
DOCQN	30.4	43.6
TF-IDF	25.3	35.4
DOC-TF-IDF	32.4	43.5
Ensemble (threshold, $l = 5$)	35.0	48.0

Table 5: Navigation accuracy for all models on the development set of TRIVIAQA-NOP.

	Development		Test	
	EM	F1	EM	F1
DQN	21.7	26.5	19.1	24.1
DOCQN	23.6	27.9	21.0	25.5
DOC-TF-IDF	21.4	27.1	18.2	23.5
Ensemble (threshold, $l = 5$)	26.8	32.0	24.2	29.4
Ensemble (answer)	28.4	33.4	25.4	30.5
READTOP	32.5	36.7	28.1	32.4

Table 6: End-to-end QA performance of all models on the development and test sets of TRIVIAQA-NOP. For the ensemble model, we choose the best threshold of $l = 5$.

	DQN	DOCQN
Path length	avg. 7.7, range [1,36]	avg. 15.2, range [3,100]
Answer predictions	avg. 2.8	avg. 4.1
Tokens consumed	3.4%	6.2%
Stopping node	0.5% title, 2.2% headline, 89.0% paragraph, 8.3% sentence	0% title, 0.4% headline, 75.1% paragraph, 24.5% sentence

Table 7: Comparing navigation properties of DQN and DOCQN on the development set. The average and range of navigation path length in steps, relative amount of consumed tokens, and distribution of stopping node types.

in sharp contrast to an ensemble with DQN, where for any value of l , DOC-TF-IDF performs better on its own.

Examining end-to-end performance, DOCQN again outperforms DQN, but DOCQN is now better than DOC-TF-IDF. This suggests that when DOC-TF-IDF selects a paragraph with the answer, it is often difficult to extract with the RC model. Again, the ensembles leads to a dramatic increase in performance, showing that DOC-TF-IDF and DOCQN are complementary. However, READTOP, which consumes the first 800 tokens of each document compared to ~ 160 for DOCQN, substantially outperforms all navigation models. This shows that the bias for answers at the beginning of documents is still strong in TRIVIAQA-NOP.

To further elucidate the differences between navigation models, Figure 7 shows navigation accuracy of different models and the proportion of samples for different node indices of the FAO. We see that DQN outperforms DOCQN when the answer is at the top of the document, but DOCQN dominates DQN when the answer is further down, showing that DOCQN learns to find answers deeper in the document. DOC-TF-IDF has a more balanced navigation accuracy across the document, which explains why an ensemble of DOC-TF-IDF with DOCQN works, as they are complementary to one another.

Analysis Figure 6 shows a navigation example, which includes the navigation step, node index, observation o , and the action a taken. For the observation we also highlight the attention distribution from the self-attention component. In this figure the question is about culture, and we see the agent going into multiple sections, reading them, running the RC model and continuing forward, until finally stopping at a paragraph that contains the answer. We provide many more examples in the supplementary material.

Table 7 highlights some differences between DQN and DOCQN. DOCQN has longer trajectories, and stops at sentence nodes more frequently, suggesting it reads the document at a finer granularity. Additionally, it leverages the RC model to collect more information and confidence during navigation, by choosing the action ANSWER more frequently. Both models consume less than 7% of the entire document. Figure 8 illustrates the navigation stopping point, and shows that DOCQN navigates deeper into the document.

6 Related Work

Handling the challenges of reasoning over multiple long documents is gaining fast momentum recently (Shen et al., 2017). As mentioned, some approaches use IR for reducing the amount of processed text (Chen et al., 2017; Clark and Gardner, 2017), while others use cheap or parallelizable models to handle long documents (Hewlett et al., 2017; Swayamdipta et al., 2018; Wang et al., 2018a). Searching for

answers while using a trained RC model as a black-box was also implemented recently in Wang et al. (2018b), for open-domain questions and multiple short evidence texts from the Web. Another thrust has focused on skimming text in a sequential manner (Yu et al., 2017), or designing recurrent architectures that can consume text quickly (Bradbury et al., 2017; Seo et al., 2018; Campos et al., 2018; Yu et al., 2018). However, to the best of our knowledge no work has previously applied these methods to long documents such as Wikipedia pages.

In this work we use TRIVIAQA-NOP for evaluation of our navigation based approach and comparison to an IR baseline. While there are various aspects to consider in such evaluation setup, our choice of data was derived mainly by the requirements for long and structured context. Recently, several new datasets such as WIKIHOP and NARRATIVEQA were published. These datasets try to focus on the tendency of RC models to match local context patterns, and are designed for multi-step reasoning. (Welbl et al., 2017; Wadhwa et al., 2018; Kočisky et al., 2017).

Our work is also related to several papers which model an agent that navigates in an environment to find objects in an image (Ba et al., 2015), relations in a knowledge-base (Das et al., 2018), or documents on the web (Nogueira and Cho, 2016).

7 Conclusions

We investigate whether document structure can be leveraged to train an agent that finds answers to questions in long documents while reading only a small fraction of it. We show that an agent that reads 6% of the document can improve QA performance compared to an IR method that utilizes the entire document, and that ensembling the two substantially improves performance. We also present DOCQN, an algorithm that promotes better exploration of the document, and show it outperforms DQN qualitatively and quantitatively.

Our approach represents a conceptual departure from previous methods for reading long documents, as it interleaves searching for an answer in the document with extracting the answer from a particular paragraph, which we show improves both navigation and QA performance. We expect that as RC models tackle longer documents that require reasoning and reading text that is spread in multiple parts of the document, models that can efficiently navigate and collect evidence will become more and more crucial. Our agent provides a first step in this important research direction.

Acknowledgments

We thank Eunsol Choi from Washington University for helpful discussions and comments on the paper. This research was partially supported by The Israel Science Foundation grant 942/16. This work was completed in partial fulfillment for the Ph.D degree of the first author.

References

- J. Ba, V. Mnih, and K. Kavukcuoglu. 2015. Multiple object recognition with visual attention. In *International Conference on Learning Representations (ICLR)*.
- J. Bradbury, S. Merity, C. Xiong, and R. Socher. 2017. Quasi-recurrent neural networks. In *International Conference on Learning Representations (ICLR)*.
- V. Campos, B. Jou, X. Giro i Nieto, J. Torres, and S. Chang. 2018. Skip RNN: Learning to skip state updates in recurrent neural networks. In *International Conference on Learning Representations (ICLR)*.
- D. Chen, J. Bolton, and C. D. Manning. 2016. A thorough examination of the CNN / Daily Mail reading comprehension task. In *Association for Computational Linguistics (ACL)*.
- D. Chen, A. Fisch, J. Weston, and A. Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Association for Computational Linguistics (ACL)*.
- J. Cheng, L. Dong, and M. Lapata. 2016. Long short-term memory-networks for machine reading. In *Empirical Methods in Natural Language Processing (EMNLP)*.

- E. Choi, D. Hewlett, A. Lacoste, I. Polosukhin, J. Uszkoreit, and J. Berant. 2017. Coarse-to-fine question answering for long documents. In *Association for Computational Linguistics (ACL)*.
- C. Clark and M. Gardner. 2017. Simple and effective multi-paragraph reading comprehension. *arXiv preprint arXiv:1710.10723*.
- R. Das, S. Dhuliawala, M. Zaheer, L. Vilnis, I. Durugkar, A. Krishnamurthy, A. Smola, and A. McCallum. 2018. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. In *International Conference on Learning Representations (ICLR)*.
- K. M. Hermann, T. Koisk, E. Grefenstette, L. Espeholt, W. Kay, M. Suleyman, and P. Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems (NIPS)*.
- M. Hessel, J. Modayil, H. V. Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver. 2017. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*.
- D. Hewlett, A. Lacoste, L. Jones, I. Polosukhin, A. Fandrianto, J. Han, M. Kelcey, and D. Berthelot. 2016. Wikireading: A novel large-scale language understanding task over Wikipedia. In *Association for Computational Linguistics (ACL)*.
- D. Hewlett, L. Jones, A. Lacoste, et al. 2017. Accurate supervised and semi-supervised machine reading for long documents. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 2011–2020.
- S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- M. Joshi, E. Choi, D. Weld, and L. Zettlemoyer. 2017. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Association for Computational Linguistics (ACL)*.
- R. Kadlec, M. Schmid, O. Bajgar, and J. Kleindienst. 2016. Text understanding with the attention sum reader network. In *Association for Computational Linguistics (ACL)*.
- T. Kočiský, J. Schwarz, P. Blunsom, C. Dyer, K. M. Hermann, G. Melis, and E. Grefenstette. 2017. The NarrativeQA reading comprehension challenge. *arXiv preprint arXiv:1712.07040*.
- N. Lao, E. Minkov, and W. Cohen. 2015. Learning relational features with backward random walks. In *Association for Computational Linguistics (ACL)*.
- K. Lee, S. Salant, T. Kwiatkowski, A. Parikh, D. Das, and J. Berant. 2017. Learning recurrent span representations for extractive question answering. *arXiv*.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- R. Nogueira and K. Cho. 2016. End-to-end goal-driven web navigation. In *Advances in Neural Information Processing Systems (NIPS)*.
- T. Onishi, H. Wang, M. Bansal, K. Gimpel, and D. McAllester. 2016. Who did what: A large-scale person-centered cloze dataset. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- S. Salant and J. Berant. 2018. Contextualized word representations for reading comprehension. In *North American Association for Computational Linguistics (NAACL)*.
- T. Schaul, J. Quan, I. Antonoglou, and D. Silver. 2015. Prioritized experience replay. In *International Conference on Learning Representations (ICLR)*.
- M. Seo, S. Min, A. Farhadi, and H. Hajishirzi. 2018. Neural speed reading via skim-RNN. In *International Conference on Learning Representations (ICLR)*.
- Y. Shen, P. Huang, J. Gao, and W. Chen. 2017. Reasonet: Learning to stop reading in machine comprehension. In *International Conference on Knowledge Discovery and Data Mining (KDD)*.
- S. Swayamdipta, A. P. Parikh, and T. Kwiatkowski. 2018. Multi-mention learning for reading comprehension with neural cascades. In *International Conference on Learning Representations (ICLR)*.

- A. Trischler, T. Wang, X. Yuan, J. Harris, A. Sordoni, P. Bachman, and K. Suleman. 2017. NewsQA: A machine comprehension dataset. In *Workshop on Representation Learning for NLP*.
- H. van Hasselt, A. Guez, and D. Silver. 2016. Deep reinforcement learning with double Q-learning. In *Association for the Advancement of Artificial Intelligence (AAAI)*, volume 16, pages 2094–2100.
- S. Wadhwa, V. Embar, M. Grabmair, and E. Nyberg. 2018. Towards inference-oriented reading comprehension: Parallelqa. In *Workshop on New Forms of Generalization in Deep Learning and Natural Language Processing at NAACL*.
- Z. Wang, T. Schaul, M. Hessel, H. V. Hasselt, M. Lanctot, and N. D. Freitas. 2016. Dueling network architectures for deep reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- S. Wang, M. Yu, X. Guo, Z. Wang, T. Klinger, W. Zhang, S. Chang, G. Tesauero, B. Zhou, and J. Jiang. 2018a. R3: Reinforced ranker-reader for open-domain question answering. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- S. Wang, M. Yu, J. Jiang, W. Zhang, X. Guo, S. Chang, Z. Wang, T. Klinger, G. Tesauero, and M. Campbell. 2018b. Evidence aggregation for answer re-ranking in open-domain question answering. In *International Conference on Learning Representations*.
- Y. Watanabe, B. Dhingra, and R. Salakhutdinov. 2017. Question answering from unstructured text by retrieval and comprehension. *arXiv preprint arXiv:1703.08885*.
- J. Welbl, P. Stenetorp, and S. Riedel. 2017. Constructing datasets for multi-hop reading comprehension across documents. *arXiv preprint arXiv:1710.06481*.
- C. Xiong, V. Zhong, and R. Socher. 2017. Dynamic coattention networks for question answering. In *International Conference on Learning Representations (ICLR)*.
- A. W. Yu, H. Lee, and Q. V. Le. 2017. Learning to skim text. In *Association for Computational Linguistics (ACL)*.
- K. Yu, Y. Liu, A. G. Schwing, and J. Peng. 2018. Fast and accurate text classification: Skimming, rereading and early stopping. In *International Conference on Learning Representations (ICLR)*.

A Supplementary Material

A.1 Network Architecture Details

Here, we elaborate on the network architecture, which was briefly described in the paper. Given an input state $s = (q, o, z, \phi_n, \phi_z)$, we denote by $q = (q_1, \dots, q_n)$, $o = (o_1, \dots, o_m)$ and $z = (z_1, \dots, z_r)$ the question tokens, observation tokens and answer prediction tokens, respectively.

Word-level embedding For every $i = 1, \dots, n$, every $j = 1, \dots, m$ and every $k = 1, \dots, r$, we create an embedding of the input tokens q_i, o_j, z_k in two steps. First, we embed every token with a pre-trained GloVe word embedding matrix W_w :

$$e_{q_i}^w = W_w \cdot q_i \quad ; \quad e_{o_j}^w = W_w \cdot o_j \quad ; \quad e_{z_k}^w = W_w \cdot z_k$$

Next, we apply character-level embeddings with a learned embedding matrix W_c . The embedded characters are then summarized with a max-pooling convolutional neural network:

$$\begin{aligned} e_{q_i}^c &= \text{CNN}(\{W_c \cdot q_{ix}\}_{x=1}^{|q_i|}) \\ e_{o_j}^c &= \text{CNN}(\{W_c \cdot o_{jx}\}_{x=1}^{|o_j|}) \\ e_{z_k}^c &= \text{CNN}(\{W_c \cdot o_{kx}\}_{x=1}^{|z_k|}) \end{aligned}$$

Concatenation of the two components yields the final word-level embeddings:

$$e_{q_i} = [e_{q_i}^w ; e_{q_i}^c] \quad ; \quad e_{o_j} = [e_{o_j}^w ; e_{o_j}^c] \quad ; \quad e_{z_k} = [e_{z_k}^w ; e_{z_k}^c]$$

Question sequence encoding The question is encoded with a BiLSTM, where for every timestamp i , the forward and backward outputs are concatenated:

$$\begin{aligned} \{u_{q_1}, \dots, u_{q_n}\} &= \text{BiLSTM}(e_{q_1}, \dots, e_{q_n}) \\ u_{q_i} &:= [\overrightarrow{\text{LSTM}}(\vec{h}_{q_{i-1}}, e_{q_i}) ; \overleftarrow{\text{LSTM}}(\overleftarrow{h}_{q_{i+1}}, e_{q_i})] \end{aligned}$$

The outputs are then summarized with a self-attention:

$$h_q = \sum_{i=1}^n \alpha_i u_{q_i}$$

where the coefficients $\alpha_1, \dots, \alpha_n$ are obtained by feeding the outputs to a two-layer feed-forward network, and normalizing the output logits with softmax:

$$a_i = \text{FFNN}(u_{q_i}) \quad ; \quad \alpha_i = \frac{e^{a_i}}{\sum_{k=1}^n e^{a_k}}$$

Observation sequence encoding The observation sequence encoding h_o is obtained in an analogous manner to h_q , except that we use a LSTM rather than a BiLSTM.

Answer prediction encoding The answer prediction encoding h_z is obtained by running the answer prediction tokens through a LSTM, and concatenating the last hidden state with the input feature vector ϕ_z :

$$\begin{aligned} \{u_{z_1}, \dots, u_{z_r}\} &= \text{LSTM}(e_{z_1}, \dots, e_{z_r}) \\ u_{z_k} &:= \overrightarrow{\text{LSTM}}(\vec{h}_{z_{k-1}}, e_{z_k}) \\ h_z &= [u_{z_r} ; \phi_z] \end{aligned}$$

	Parameter	Value
Network	Maximum node observation length	20 tokens
	Maximum observation length	120 tokens
	Word embedding dimension	300
	Character embedding dimension	20
	Convolution filter size	5
	BiLSTM and LSTM hidden dimension	300
	First feed-forward layer dimension	512
	Second feed-forward layer dimension	256
	Dropout rate	0.2
	Training	RMSprop learning rate
Batch size		64
Target network period		10K steps
Initial memory size		50K transitions
Maximal memory size		300K transitions
Action sampling ϵ		1.0 \rightarrow 0.1
Discount factor γ		0.996
Prioritization usage α		0.6
Prioritization importance sampling β		0.4 \rightarrow 1.0
Sampling	State sampling ϵ_s	1.0 \rightarrow 0.5
	Annealing steps for ϵ_s	1.2M steps
	Sampling repetitions K	5
	Maximum navigation length T (train)	30 steps
	Maximum navigation length T (test)	100 steps
	Interpolation coefficient for f_{U+B}	0.5

Table 8: DOCQN hyper-parameters

State representation The final state representation is formed by concatenating the encoded question h_q with the encoded observation h_o . Concretely:

$$h_s = [h_q ; h_o ; h_z]$$

The input node features ϕ_o are concatenated to an upper layer, as we describe next, to increase their weights in the final predictions. We have found that incorporation of these features in this way accelerates the navigation learning process.

Q-values prediction We implement a Dueling DQN architecture, where the final Q-values are composed of a state value $V(s)$ prediction and advantage predictions $A(s, a)$ for every possible action a . Denoting by FFNN a single-layer feed-forward neural network, predictions are derived as follows:

$$\begin{aligned} v_0 &= \text{FFNN}(h_s) \\ v_1^V &= \text{FFNN}(v_0) ; v_1^A = \text{FFNN}(v_0) \\ v_2^V &= \text{FFNN}([v_1^V ; \phi_n]) ; v_2^A = \text{FFNN}([v_1^A ; \phi_n]) \end{aligned}$$

Where $v_2^V \in \mathbb{R}$, $v_2^A \in \mathbb{R}^{|\mathcal{A}|}$, and ϕ are the navigation features. The final Q-values prediction is obtained by averaging:

$$v^Q = v_2^V + \left(v_2^A - \frac{\sum_{i=1}^{|\mathcal{A}|} v_{2i}^A}{|\mathcal{A}|} \right)$$

A.2 Hyper Parameters

Table 8 summarizes the hyper-parameters used for building and training the DOCQN models.

A.3 Navigation Examples

Figures 9,10,11,12,13 show sample navigations, performed by the DOCQN model. In each example, q, a, n denote the question, answer aliases and answer node numbers. The observation tokens are highlighted according to the self-attention weights, given by the model. By choosing the action ANS, the agent executes the RC model to obtain an answer prediction, which is part of the observation in the next step.

q: How many legs does a ladybird have?
a: [six', '6']
n: [2, 12, 14, 22, 46, 51, 60, 72, 74, 102, 150, 159, 163, 200]

Step	Node	Observation Tokens	Answer	Action
0	0	['Insect']		DOWN
1	1	['Insect', 'Etymology']		DOWN
2	2	['Insect', 'Etymology', 'The', 'word', 'is', 'insect', 'comes', 'from', 'the', 'Latin', 'word', 'is', 'meaning', 'is', 'with', 'a', 'notched', 'or', 'divided', 'body']		UPR
3	5	['Insect', 'Phylogeny', 'and', 'evolution']		DOWN
4	6	['Insect', 'Phylogeny', 'and', 'evolution', 'The', 'evolutionary', 'relationship', 'of', 'insects', 'to', 'other', 'animal', 'groups', 'remains', 'unclear', '']		UPR
5	45	['Insect', 'Diversity']		DOWN
6	46	['Insect', 'Diversity', 'Thought', 'the', 'true', 'dimensions', 'of', 'species', 'diversity', 'remain', 'uncertain', 'estimates', 'range', 'from', '2.6', '7.8', 'million', 'species', 'with', 'a']		UPR
7	49	['Insect', 'Morphology', 'and', 'physiology']		DOWN
8	50	['Insect', 'Morphology', 'and', 'physiology', 'External']		DOWN
9	51	['Insect', 'Morphology', 'and', 'physiology', 'External', 'Insects', 'have', 'segmented', 'bodies', 'supported', 'by', 'exoskeletons', 'the', 'hard', 'outer', 'covering', 'made', 'mostly', 'of', 'chitin', '']		ANS
10	51	['Insect', 'Morphology', 'and', 'physiology', 'External', 'Insects', 'have', 'segmented', 'bodies', 'supported', 'by', 'exoskeletons', 'the', 'hard', 'outer', 'covering', 'made', 'mostly', 'of', 'chitin', '']	six	STOP

Figure 9: Navigation example 1. Although the document is more general than the question subject and the answer appears multiple times across the document, the agent finds the correct context for answering the question.

q: What was the French sounding winner of the 2011 Epsom Derby?
a: [pour moi']
n: [15, 40]

Step	Node	Observation Tokens	Answer	Action
0	0	['2011', 'Epsom', 'Derby']		DOWN
1	1	['2011', 'Epsom', 'Derby', 'Race', 'details']		DOWN
2	2	['2011', 'Epsom', 'Derby', 'Race', 'details', 'Sponsor', '']		UPR
3	9	['2011', 'Epsom', 'Derby', 'Full', 'result']		DOWN
4	10	['2011', 'Epsom', 'Derby', 'Full', 'result', 'The', 'distances', 'between', 'the', 'horses', 'are', 'shown', 'in', 'lengths', 'or', 'shorter', '']		UPR
5	14	['2011', 'Epsom', 'Derby', 'Winner', 'details']		DOWN
6	15	['2011', 'Epsom', 'Derby', 'Winner', 'details', 'Further', 'details', 'of', 'the', 'winner', 'Pour', 'Moi', '']		STOP

Figure 11: Navigation example 3. The agent quickly navigates to the answer, without running the RC model.

q: Justine Thornton is the fiancée of which politician?
a: [ed milliband', 'red ed', 'tom baldwin journalist', 'ed miliband', 'edward miliband']
n: [21]

Step	Node	Observation Tokens	Answer	Action
0	0	['Justine', 'Thornton']		DOWN
1	1	['Justine', 'Thornton', 'Early', 'life', 'and', 'education']		DOWN
2	2	['Justine', 'Thornton', 'Early', 'life', 'and', 'education', 'Thornton', 'was', 'born', 'in', 'Manchester', 'to', 'professionals', 'Margaret', 'and', 'Dr', 'Stewart', 'Thornton', '']		ANS
3	2	['Justine', 'Thornton', 'Early', 'life', 'and', 'education', 'Thornton', 'was', 'born', 'in', 'Manchester', 'to', 'professionals', 'Margaret', 'and', 'Dr', 'Stewart', 'Thornton', '']	Dr Stewart	UPR
4	10	['Justine', 'Thornton', 'Career']		DOWN
5	11	['Justine', 'Thornton', 'Career', 'Thornton', 'practises', 'in', 'Environmental', 'Law', 'now', 'at', '39', 'Essex', 'Street', '']		ANS
6	11	['Justine', 'Thornton', 'Career', 'Thornton', 'practises', 'in', 'Environmental', 'Law', 'now', 'at', '39', 'Essex', 'Street', '']	Environmental Law	UPR
7	20	['Justine', 'Thornton', 'Personal', 'life']		DOWN
8	21	['Justine', 'Thornton', 'Personal', 'life', 'Thornton', 'is', 'husband', 'is', 'former', 'Labour', 'Party', 'Leader', 'Ed', 'Miliband', 'MP', '']		STOP

Figure 10: Navigation example 2. The agent explores several sections before reaching the most relevant one.

q: What nationality was the pointillist artist MONDRIAN?
a: [autochtoon', 'people of holland', 'ethnic dutch', 'dutch person', 'dutch', 'dutch people', 'dutch ethnic group', 'ethnic dutch', 'dutch ethnic group and nation', 'netherlanders', 'autochtones', 'dutchmen', 'autochtonen', 'netherlanders', 'netherlander']
n: [3, 9, 53]

Step	Node	Observation Tokens	Answer	Action
0	0	['Piet', 'Mondrian']		DOWN
1	1	['Piet', 'Mondrian', 'The', 'Netherlands', '1872', '1911', '']		DOWN
2	2	['Piet', 'Mondrian', 'The', 'Netherlands', '1872', '1911', 'Mondrian', 'was', 'born', 'in', 'in', 'the', 'Netherlands', 'the', 'second', 'of', 'his', 'parents', 'children', '']		ANS
3	2	['Piet', 'Mondrian', 'The', 'Netherlands', '1872', '1911', 'Mondrian', 'was', 'born', 'in', 'in', 'the', 'Netherlands', 'the', 'second', 'of', 'his', 'parents', 'children', '']	Netherlands	RIGHT
4	3	['Piet', 'Mondrian', 'The', 'Netherlands', '1872', '1911', 'Mondrian', 'was', 'born', 'in', 'in', 'the', 'Netherlands', 'the', 'second', 'of', 'his', 'parents', 'children', '']		ANS
5	3	['Piet', 'Mondrian', 'The', 'Netherlands', '1872', '1911', 'Mondrian', 'was', 'born', 'in', 'in', 'the', 'Netherlands', 'the', 'second', 'of', 'his', 'parents', 'children', '']	Dutch	STOP

Figure 12: Navigation example 4. The agent leverages the RC model and decides to stop after observing enough information to answer.

q: Who created Rumpole of the Bailey?
a: [john mortimer qc', 'john mortimer', 'john clifford mortimer', 'sir john mortimer']
n: [31, 46, 56, 144, 154, 168, 178, 188, 198, 208, 245, 258, 274, 282, 290, 306, 341, 345, 466, 475]

Step	Node	Observation Tokens	Answer	Action
0	0	['Rumpole', 'of', 'the', 'Bailey']		DOWN
1	1	['Rumpole', 'of', 'the', 'Bailey', 'Horace', 'Rumpole']		RIGHT
2	44	['Rumpole', 'of', 'the', 'Bailey', 'Production']		RIGHT
3	127	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes']		DOWN
4	128	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', 'There', 'were', 'a', 'grand', 'total', 'of', '44', 'episodes', '']]		RIGHT
5	129	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', 'All', 'listed', 'dates', 'indicate', 'first', 'UK', 'transmission', 'date']		RIGHT
6	130	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', ']', 'Film', 'for', 'BBC', 'TV', 's', 'Play', 'for', 'Today', 'Series', '(, '1975', ')]]		RIGHT
7	131	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', '*', ']', 'Rumpole', 'of', 'the', 'Bailey', 'w', '(, '16', 'December', '1975', ')], T, 'Set', 'in', '1974', T]		RIGHT
8	132	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', '**', '(, ', ']', 'Rumpole', 'and', 'the', 'Confession', 'of', 'Guilt', 'w', 'for', 'radio', 'adaptation', 'in', '1980', 'and', 'for', 'DVD', 'release']		RIGHT
9	133	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', '*', 'This', 'was', 'a', 'stand-alone', 'production', 'in', '1975', 'for', 'BBC', 'TV', 's', 'anthology', 'series', 'w', 'Play', 'for', 'Today', 'w', '']]		RIGHT
10	134	['Rumpole', 'of', 'the', 'Bailey', 'Television', 'episodes', 'TV', 'Season', 'One', '(, '1978', ')]]		UPL
11	44	['Rumpole', 'of', 'the', 'Bailey', 'Production']		DOWN
12	45	['Rumpole', 'of', 'the', 'Bailey', 'Production', 'Origins']		DOWN
13	46	['Rumpole', 'of', 'the', 'Bailey', 'Production', 'Origins', 'The', 'origins', 'of', 'Rumpole', 'of', 'the', 'Bailey', 'lie', 'in', 'w', ']', 'Took', 'Place', 'w', ']', 'a', 'one-off', 'filmed', 'television', 'play']		ANS
14	46	['Rumpole', 'of', 'the', 'Bailey', 'Production', 'Origins', 'The', 'origins', 'of', 'Rumpole', 'of', 'the', 'Bailey', 'lie', 'in', 'w', ']', 'Took', 'Place', 'w', ']', 'a', 'one-off', 'filmed', 'television', 'play']	John Mortimer	STOP

Figure 13: Navigation example 5. After reading through the "Television episodes" section, the agent goes back to the "Production" section to find the answer.