

Article

ClueReader: Heterogeneous Graph Attention Network for Multi-Hop Machine Reading Comprehension

Peng Gao ^{1,2,*} , Feng Gao ^{1,3}, Peng Wang ¹, Jian-Cheng Ni ⁴, Fei Wang ⁵ and Hamido Fujita ^{6,7,8} ¹ School of Cyber Science and Engineering, Qufu Normal University, Qufu 273165, China² Yuntian Group, Qingyun 253700, China³ School of Computer Science and Technology, East China Normal University, Shanghai 200062, China⁴ Network and Information Center, Qufu Normal University, Qufu 273165, China⁵ School of Electronics and Information Engineering, Harbin Institute of Technology, Shenzhen 518055, China⁶ Faculty of Information Technology, HUTECH University, Ho Chi Minh City 70000, Vietnam⁷ Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI), University of Granada, 18011 Granada, Spain⁸ Research and Regional Cooperation Division, Iwate Prefectural University, Iwate 020-0611, Japan

* Correspondence: pgao@qfnu.edu.cn

Abstract: Multi-hop machine reading comprehension is a challenging task in natural language processing as it requires more reasoning ability across multiple documents. Spectral models based on graph convolutional networks have shown good inferring abilities and lead to competitive results. However, the analysis and reasoning of some are inconsistent with those of humans. Inspired by the concept of *grandmother cells* in cognitive neuroscience, we propose a heterogeneous graph attention network model named *ClueReader* to imitate the *grandmother cell* concept. The model is designed to assemble the semantic features in multi-level representations and automatically concentrate or alleviate information for reasoning through the attention mechanism. The name *ClueReader* is a metaphor for the pattern of the model: it regards the subjects of queries as the starting points of clues, takes the reasoning entities as bridge points, considers the latent candidate entities as *grandmother cells*, and the clues end up in candidate entities. The proposed model enables the visualization of the reasoning graph, making it possible to analyze the importance of edges connecting entities and the selectivity in the mention and candidate nodes, which is easier to comprehend empirically. Evaluations on the open-domain multi-hop reading dataset WIKIHOP and drug–drug interaction dataset MEDHOP proved the validity of *ClueReader* and showed the feasibility of its application of the model in the molecular biology domain.

Keywords: machine reading comprehension; knowledge graph; graph neural networks; attention mechanism



Citation: Gao, P.; Gao, F.; Wang, P.; Ni, J.-C.; Wang, F.; Fujita, H.

ClueReader: Heterogeneous Graph Attention Network for Multi-Hop Machine Reading Comprehension. *Electronics* **2023**, *12*, 3183. <https://doi.org/10.3390/electronics12143183>

Academic Editor: Rui Pedro Lopes

Received: 22 June 2023

Revised: 11 July 2023

Accepted: 21 July 2023

Published: 22 July 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Machine reading comprehension (MRC) is one of the most attractive and long-standing tasks in natural language processing (NLP). Compared with single-paragraph MRC, multi-hop MRC is more challenging since multiple confusing answer candidates are contained in different passages [1,2]. Models designed for this task are supposed to have abilities to reasonably traverse multiple passages and discover reasoning clues following given questions. For complex multi-hop MRC tasks, more understandable, reliable, and analyzable methodologies are required to improve reading performance.

Better understanding of biological brains could play a vital role in building artificial intelligent systems [3]. Previous cognitive research in reading can be of benefit to challenging multi-hop MRC tasks. The concept of *grandmother cells* can be traced back to a 1969 academic lecture given by the neuroscientist Jerome Lettvin [4] and was later defined by the physiologist Horace Barlow as cells in the brain that respond specifically to a single

familiar person or object. In experiments on primates, researchers discovered individual neurons that responded specifically to a specific person, image, or concept after differentiation [5]. A study of a patient with epilepsy found a neuron in the patient's anterior temporal lobe that responded specifically to the Hollywood star Jennifer Aniston [6]. Any form of stimulation related to Aniston, whether it be a color photograph, a close-up of her face, or a cartoon portrait, or even just seeing her name written on paper, could and would only stimulate that neuron to produce an excited signal. As research into the concept of *grandmother cells*, the underlying mechanism of their response became clearer. The signal output from a single *grandmother cell* in response to specific stimuli actually stems from the coordinated calculation of a large-scale neural network behind *grandmother cells* [5]. It suggests that a single neuron can respond to only one out of thousands of stimulations, which is somehow intuitively similar to reading and inference in multi-hop MRC:

- **Selectivity.** The *grandmother cells* concept organizes the neurons in a hierarchical "sparse" coding scheme. It activates some specific neurons to respond to stimulation, similar to the manner in which we store reasoning evidence maps (neurons) in our minds during reading and recall-related evidence maps to reason the answer with a question (stimulation) constrained.
- **Specificity.** The concept implies that brains contain *grandmother* neurons that are so specialized and dedicated to a specific object, which is similar to a particular MRC question resulting in a specific answer among multiple reading passages and their complex reasoning evidence.
- **Class character.** Amazing selectivity is captured in *grandmother cells*. However, it results from computation by much larger networks and the collective operations of many functionally different low-level cells, similar to human multi-hop reading in which evidence is usually gathered from different levels as much as possible and the final answer is decided in some candidate endpoints.

To imitate *grandmother cells* in multi-hop MRC, the reading evidence is supposed to be organized as level-classified neurons and the selections must be performed in response to specific question stimulation. As for multi-hop MRC tasks, the hops between two entities could be connected as node pairs and gradually constructed into a reasoning evidence graph taking all related entities as nodes. This reasoning evidence graph is intuitively represented as a graph structure, which can be empirically considered to contain the implicit reasoning chains from the start of the question to the end of the answer nodes (entities). We generally recall considerable related evidence as a node, whatever form it is (such as a paragraph, a short sentence, or a phrase), to meet the class character, and we coordinate their inter-relationship before obtaining the results.

Graph neural networks (GNNs) inspire us to posit that operating on graphs and manipulating the structured knowledge can support relational reasoning [7,8] in a sophisticated and flexible pattern, similar to the implementation of *grandmother cells* regarding the cells as nodes in the graph and collecting evidence in multi-classified aspects of node representations. Further, spatial graph attention networks (GATs) perform the selectivity in reasoning evidence graphs in the manner of *grandmother cells* using attention mechanisms. This work has the following main contributions:

1. In order to construct a more reasonable graph, *ClueReader* draws inspiration from the concept of *grandmother cells* in the brain during information cognition, in which cells in the brain only output specific entities. This leads to the creation of heterogeneous graph attention networks with multiple types of nodes.
2. By taking the subject of queries as the starting point, potential reasoning entities in multiple documents as bridge points, and mention entities consistent with candidate answers as end points, the proposed *ClueReader* is a heuristic way of constructing MRC chains.
3. Before outputting predicted answers, *ClueReader* innovatively visualizes the internal state of the heterogeneous graph attention network, providing intuitive quantitative data displays for analyzing the effectiveness, rationality, and explainability.

The remainder of the article is organized as follows. Section 2 describes the work related to multi-hop MRC, and Section 3 proposes the *ClueReader* that imitates *grandmother cells* for multi-hop MRC. Experimental evaluations are conducted in Section 4, and conclusions are summarized in Section 5.

2. Related Work

2.1. Sequential Reading Models for Multi-Hop MRC

Sequential reading models were first used for single-passage MRC tasks, and most of them are based on recurrent neural networks (RNNs) or their variants. When the attention mechanism was introduced into NLP tasks, their performance significantly improved [9–12]. In the initial benchmarks of the QANGAROO [13], a dataset for multi-hop MRC, the milestone model *Bi-Directional Attention Flow (BiDAF)* [9] was first applied to evaluate its performance in the multi-hop MRC task. It represented the context at different levels and used a bi-directional attention flow mechanism to obtain query-aware context representation and was then used for predictions.

Some studies [14–17] argued that independent attention mechanisms, i.e., *Bidirectional Encoder Representations from Transformers (BERT)*-style models [14], applied on sequential contexts can outperform former RNN-based approaches in various NLP downstream tasks, including MRC. When the sequential approaches were applied to multi-hop MRC tasks, however, they suffered from the challenge that the super-long contexts—to adapt the design of the sequential requirement, multiple passages are concatenated into one passage—resulted in dramatically increased calculation and time consumption. A long-sequence architecture, *Longformer* [17], overcomes the self-attention restriction and allows the length of sequences to be increased from 512 to 4096 and then concatenates all the passages into a long sequential context for reading. The *Longformer* modified the question answering (QA) methodology proposed in *BERT* [14]: the long sequential context consisted of a question, candidates, and passages, which were separated by special tags that were applied to the linear layers to output the predictions, while still having enough memory for first 4096 length sequence.

Although the approaches above are effective, Ref. [18] indicates that model reasoning is not robust enough. We consider that there are still two main challenges that should be further addressed: (1) With the expansion of the problem scale and the reasoning complexity, the token-limited problem may appear again eventually. For instance, a full-wiki setting task in HOTPOTQA requires models to predict answers from the scope of the entire WIKIPEDIA, which is a dataset for diverse and explainable multi-hop question answering. It is difficult to imagine how a huge search space is built based on a large amount of text. (2) Some models which simply concatenate text to long contexts lack logical relationships, which is unconvincing in terms of their reasoning. Thus, the approaches based on GNNs were proposed to improve the scalability and explainability in multi-hop MRC.

2.2. Graph Neural Networks for Multi-Hop MRC

Reasoning about explicitly structured data, in particular, graphs, has arisen at the intersection of deep learning and structured approaches [7]. As the representative graph methodology, *Graph Convolutional Networks (GCNs)* [19,20] are widely applied in multi-hop MRC approaches. *Cognitive Graph QA (CogQA)* [21] was founded on the dual process theory [22,23], and it divides the multi-hop reading process into two stages: the implicit extraction (System I) based on *BERT* and the explicit reasoning (System II) established in GCNs. System I extracts the answer candidates and useful next-hop entities from passages for the cognitive graph construction, then System II updates entity representations and predicts the final answer in the GCN message passing way. In this procedure, the selected passages are not put in the system at once. As a result, *CogQA* keeps its scalability in the face of the massive scope of reading materials.

Entity-GCN [24] extracts all the text spans matching the candidates as nodes and obtains their representations from the contextualized *ELMo* [25] word embeddings, then passes them to the GCN module for reasoning. Based on *Entity-GCN*, *Bi-directional Attention Entity Graph Convolutional Network (BAG)* [26] added *Glove* word embeddings and two manual features, named-entity recognition and part-of-speech tags, to reflect the semantic properties of tokens. On account of the full usage of the question contextual information, it applies the bi-directional attention mechanism, both node2query and query2node, to obtain query-aware node representations in the reasoning graph for better predictions. *Path-based GCN* [27] introduces more related entities in the graph than the nodes merely matching the candidates to enhance the performance of the model. The *Heterogeneous Document-Entity (HDE)* model [28] introduces the heterogeneous nodes into GCNs, which contain different granularity levels of information. Additionally, the *Keywords-Aware Dynamic Graph Neural Network (KA-DGN)* [29] was proposed and designed as a dynamic graph neural network to further tackle reading over multiple scattered text snippets. Furthermore, Zhan et al. [30] and Song et al. [31] separately proposed knowledge-aware and evidence-aware GNN reading models, which integrate dependency relations or multiple pieces of evidence from multiple paragraphs.

However, the reading processes of the above-mentioned approaches are still inexplicable, especially in GNNs, which stimulated our interest in the selectivity of this procedure.

3. Methodology

We introduce the design and implementation of the proposed model, *ClueReader*, which is shown in Figure 1.

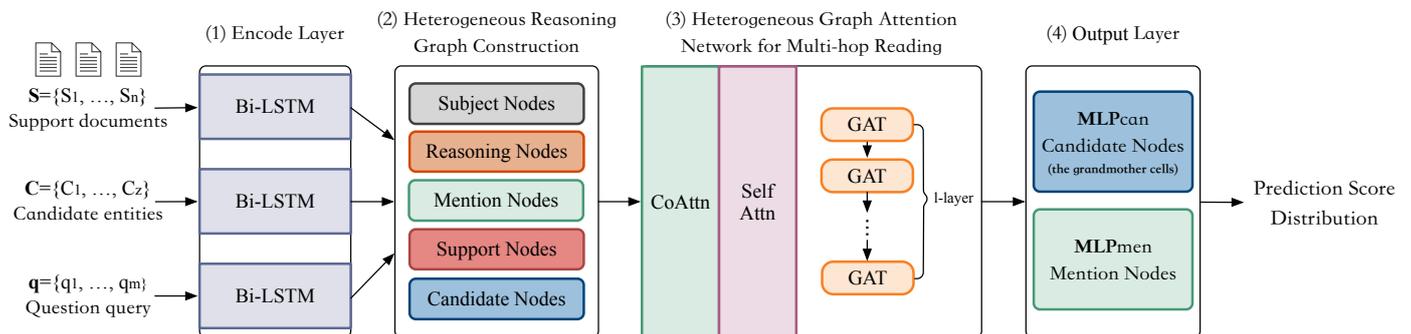


Figure 1. Our proposed *ClueReader*: a heterogeneous graph attention network for multi-hop MRC. The detailed explanations of S , C , and q are in task formalization (Section 3.1). S , C , and q are encoded in three independent *Bi-LSTM*s (Section 3.2). Following the graph construction strategies in Section 3.3, the outputs of three encoders are applied to *Co-attention* and *Self-attention* to initialize the reasoning graph features, which is explained in Section 3.4. Then the topology information and node features are passed into the GAT layer. A much larger network computation behind *grandmother cells* is performed in the GAT layer, and n -hops message passing is calculated in n parameter shared layers, which are represented in Section 3.4.2. Finally, *grandmother cell* selectivity is combined in Section 3.5, outputting the final predicted answer.

3.1. Task Formalization

A given query $q = (s, r, a^*)$ is in a triple form, where s is the subject entity, r is the query relation (i.e., predication), and q can be converted into sequential form $q = \{q_1, q_2, \dots, q_m\}$, where m is the number of tokens in the query q . Then a set of candidates $C_q = \{c_1, c_2, \dots, c_z\}$ and a series of supporting documents $S_q = \{s_1, s_2, \dots, s_n\}$ containing the candidates are also provided, where z is the number of the given candidates, n is the number of the given supporting documents, and the subscript q means the two sets are constrained by the query q . Moreover, S_q is provided in a random order, and without S_q , the answer to the query q could be multiple. Our goal is to identify the single correct answer $a^* \in C_q$ by reading S_q .

3.2. Encoding Layer

We utilize the pre-trained *GloVe* [32] model to initialize word embeddings, and then employ *Bidirectional Long Short-Term Memory (Bi-LSTM)* [33,34] to encode sequence representations as:

$$\begin{aligned} \begin{pmatrix} f_t \\ i_t \\ o_t \end{pmatrix} &= \sigma(W_h h_{t-1} + W_i x_t) \\ \tilde{c}_t &= \tanh(W_h h_{t-1} + W_i x_t) \\ c_t &= f_t c_{t-1} + i_t \tilde{c}_t \\ h_t &= o_t \tanh(c_t) \end{aligned} \quad (1)$$

where the subscripts t and $t - 1$ denote the indexes of the encoding time step; W_i and W_h are the hyperparameters of the input and the hidden layer; i , f , o , \tilde{c} , h , and c respectively represent the input, forget, output, content, hidden, and cell states; x represents the word embedding; σ and \tanh are sigmoid activation and hyperbolic tangent activation, respectively.

We use \vec{h} and \overleftarrow{h} to denote the forward-pass (i.e., the left to right) and the backward-pass (i.e., the right to left) sequence representations encoded by *Bi-LSTM*, respectively. Then, the representation of the entire sequential context obtained from the encoding layer can be expressed as follows:

$$h = [\vec{h} || \overleftarrow{h}] \quad (2)$$

where the symbol $||$ denotes the concatenation of \vec{h} and \overleftarrow{h} . To encode the sequence representations of support documents S , candidates C , and query q , it is desirable to use three independent *Bi-LSTMs*. Their outputs are $H_s^i \in \mathbb{R}^{l_s \times d}$, $H_c^j \in \mathbb{R}^{l_c \times d}$ and $H_q \in \mathbb{R}^{l_q \times d}$, respectively, where i and j are the indexes of the documents and the candidates, l is the sequence length, and d is the output dimension of the representations.

3.3. Heterogeneous Reasoning Graph

The concept of *grandmother cells* reveals that the brains of monkeys, like those of humans, contain neurons that are so specialized they appear to be dedicated to a single person, image, or concept. This amazing selectivity is uncovered in a single neuron, while it must result from computation by a much larger network [5]. We heuristically consider that this procedure in multi-hop reading could be summarized as three steps:

1. The query (or the question) locates the related neurons at a low level, which then stimulates higher-level neurons to trigger computation;
2. The higher-level neurons begin to respond to increasingly broader portions of other neurons for reasoning, and to avoid a broadcast storm, informative selectivity takes place in this step;
3. At the top-level, some independent neurons are responsible for the computations that occurred in step 2. We refer to these neurons as *grandmother cells* and expect them to provide the appropriate results that correspond to the query.

We attempt to imitate *grandmother cells* in our reading procedure and present our reasoning graph as consistent as possible with the three steps mentioned above. The heterogeneous reasoning graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, which is illustrated in Figure 2, simulates a heuristic chain of comprehension that starts from the subject entity in query q and goes through the reasoning entities in the supporting document set S_q , then through the mention entities in S_q that are consistent with the candidate answer, and finally touches at the candidates in set C_q (referred to as the *grandmother cell*).

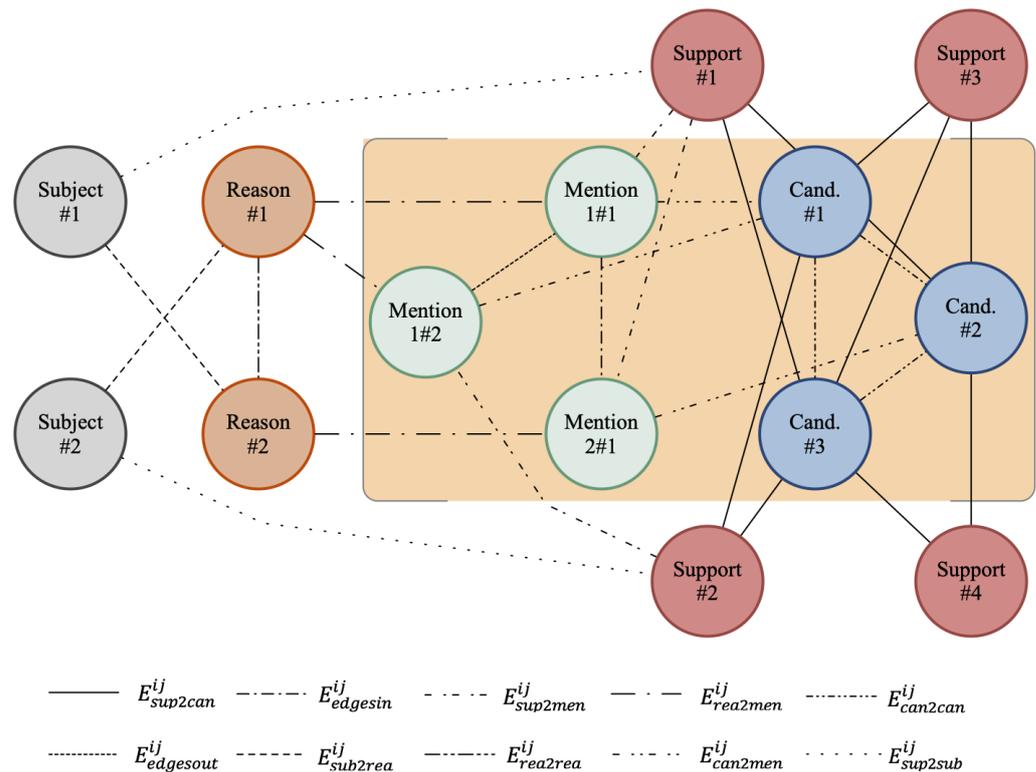


Figure 2. Heterogeneous reasoning graph in *ClueReader*. Different nodes are filled in different colors, and the edges are distinguished by the types of lines. Subject nodes are gray, reasoning nodes are orange, mention nodes are green, support nodes are red, and candidate nodes are blue. The nodes in the light yellow square are all selected to input to the two MLP obtaining the prediction score distribution.

3.3.1. Nodes Definition

To construct the graph, we define five different types of nodes which are similar to neurons and ten kinds of edges among the nodes [15,24].

- **Subject Nodes**—As the form of query q , the subject entity s is given in $q = (s, r, a^*)$. For example, the subject entity of the query sequence context *Where is the basketball team that Mike DiNunno plays for based?* is certainly *Mike DiNuuno*. We extract all the named entities that match with s from documents, and regard them as the subject nodes to open up the reading clues triggering the further computations. The subject nodes are denoted as \mathcal{V}_{sub} and colored in gray in Figure 2.
- **Reasoning Nodes**—In light of the requirements of the multi-hop MRC, there are some gaps between the subject entities and candidates. To build bridges between the two and make the reasoning clues as complete as possible, we replenish those clues with the named recognition entities and nominal phrases from the documents containing the question subjects and answer candidates. The reasoning nodes are marked as \mathcal{V}_{rea} and colored in orange in Figure 2.
- **Mention Nodes**—A series of candidate entities are given in C_q , they may occur in multiple times within the document set S_q . As a result, we traverse the documents and extract the named entities corresponding to each candidate as mention nodes, serving as the soft endpoint of the reasoning chain. It should be noted that mention nodes will participate in the semi-supervised learning process and will be involved in the final answer prediction. The mention nodes are presented as \mathcal{V}_{men} and colored in green in Figure 2.
- **Support Nodes**—As described by [5], we consider that multi-type representations may contribute to the reading process, and thus the support documents containing

the above nodes are introduced to \mathcal{G} as support nodes, which are notated as \mathcal{V}_{sup} and colored in red in Figure 2.

- **Candidate Nodes**—To imitate *grandmother cells*, we consider candidate nodes as hard endpoints of the reasoning chain to gather relevant information from the heterogeneous reasoning graph. For the mention nodes \mathcal{V}_{men}^q of a candidate answer c_q , when $\mathcal{V}_{men}^q \geq 1$, candidate nodes are established as *grandmother cells* to provide the final prediction. The candidate nodes are denoted as \mathcal{V}_{can} and colored in blue in Figure 2.

3.3.2. Edges Definition

To learn the entity relationships between different nodes, we define 10 kinds of edges between nodes in heterogeneous reasoning graphs inspired by the literature [24,26,35], as shown in Table 1.

Table 1. The definition of edges in the heterogeneous graph attention network *ClueReader*.

Edges	Definition
$\mathcal{E}_{sup2sub}$	If the support document s_i contains the j -th subject node v_{sub}^j , an undirected edge denoted as $e_{sup2sub}^{ij}$ is established to connect the support node v_{sup}^i of s_i and the subject node v_{sub}^j .
$\mathcal{E}_{sup2can}$	If the support document s_i contains the j -th candidate node v_{can}^j , an undirected edge denoted as $e_{sup2can}^{ij}$ is established to connect the support node v_{sup}^i of s_i and the candidate node v_{can}^j .
$\mathcal{E}_{sup2men}$	If the support document s_i contains the j -th mention node v_{men}^j , an undirected edge denoted as $e_{sup2men}^{ij}$ is established to connect the support node v_{sup}^i of s_i and the mention node v_{men}^j .
$\mathcal{E}_{can2men}$	If the j -th mention node v_{men}^j and the i -th candidate node v_{can}^i represent the same entity, an undirected edge denoted as $e_{can2men}^{ij}$ is established to connect the two nodes.
$\mathcal{E}_{sub2rea}$	If the i -th subject node v_{sub}^i and the j -th reasoning node v_{rea}^j extracted from the same document, an undirected edge denoted as $e_{sub2rea}^{ij}$ is established to connect the two nodes.
$\mathcal{E}_{rea2men}$	If the i -th reasoning node v_{rea}^i and the j -th mention node v_{men}^j extracted from the same document, an undirected edge denoted as $e_{rea2men}^{ij}$ is established to connect the two nodes.
$\mathcal{E}_{can2can}$	All the mention nodes are fully connected using undirected edge $e_{can2can}^{ij}$.
$\mathcal{E}_{edgesin}$	If two mention nodes v_{men}^i and v_{men}^j are extracted from the same document, the two nodes will be connected as $e_{edgesin}^{ij}$.
$\mathcal{E}_{edgesout}$	If two mention nodes v_{men}^i and v_{men}^j are extracted from different documents represent the same entity, the two nodes will be connected as $e_{edgesout}^{ij}$.
$\mathcal{E}_{rea2rea}$	If two reasoning nodes v_{rea}^i and v_{rea}^j are extracted from the same document or represent the same entity, the two nodes will be connected as $e_{rea2rea}^{ij}$.

3.3.3. Graph Construction

In the heterogeneous reasoning graph, the clue-reading chain can be represented by $\mathcal{V}_{sub} \leftrightarrow \mathcal{V}_{rea} \leftrightarrow \mathcal{V}_{men} \leftrightarrow \mathcal{V}_{can}$, whose edges are covered by $\mathcal{E}_{sub2rea}$, $\mathcal{E}_{rea2rea}$, $\mathcal{E}_{rea2men}$, and $\mathcal{E}_{can2men}$. $\mathcal{E}_{edgesout}$ and $\mathcal{E}_{rea2rea}$ give the model the ability to transfer information across documents and edges in $\mathcal{E}_{sup2sub}$, $\mathcal{E}_{sup2can}$, and $\mathcal{E}_{sup2men}$ are responsible for supplementing the multi-angle textual information from the documents. Furthermore, the $\mathcal{E}_{can2men}$ could gather all the information of the mentioned nodes corresponding to the candidates and then pass their representations to the output layer to realize the imitation of *grandmother cells*.

Specifically, this multi-hop MRC process of the clue-based reasoning starts with the subject node, connecting reasoning nodes from support documents, then connecting the mention nodes as soft endpoints of the clue chain, and finally connecting to the candidate nodes (*grandmother cells*) as hard endpoints of the clue chain. For example, for the question *Which country is the location of the United Nations Headquarters?* the answer candidate set

includes *China, France, UK, USA, and Russia*. One correct and reasonable clue chain can be represented as *Location of United Nations Headquarters* (subject node) ↔ *Manhattan* ↔ *New York City* ↔ *New York State* ↔ *USA* (mention node) ↔ *USA* (candidate node). In practice, multiple clue chains are included within the heterogeneous reasoning graph, and under the constraints of the query, the selection of soft and hard endpoints is required to output the final prediction.

3.4. Heterogeneous Graph Attention Network for Multi-Hop Reading

3.4.1. Query-Aware Contextual Information

Following *HDE* [28], we use the *co-attention* and *self-attention* mechanisms [36] to combine the query contextual information and documents. Moreover, it is applied to the other semantic representations that require reasoning consistent with the query. To represent the query-aware support documents, it can be calculated as follows:

$$\mathcal{A}_{qs}^i = H_s^i (H_q)^\top \in \mathbb{R}^{l_s \times l_q} \tag{3}$$

where \mathcal{A}_{qs}^i is the similarity matrix for two sequences, between the i -th support document $H_s^i \in \mathbb{R}^{l_s \times d}$ and query $H_q \in \mathbb{R}^{l_q \times d}$, and d is the dimension of the context. Then, the query-aware representation of support documents S_{ca} is computed as follows:

$$\mathcal{K}_q = \text{softmax}(\mathcal{A}_{qs}^\top) H_s \in \mathbb{R}^{l_q \times d} \tag{4}$$

$$\mathcal{K}_s = \text{softmax}(\mathcal{A}_{qs}) H_q \in \mathbb{R}^{l_s \times d} \tag{5}$$

$$\mathcal{D}_s = \text{BiLSTM}(\text{softmax}(\mathcal{A}_{qs}) \mathcal{K}_q) \in \mathbb{R}^{l_s \times d} \tag{6}$$

$$S_{ca} = [\mathcal{K}_s || \mathcal{D}_s] \in \mathbb{R}^{l_s \times 2d} \tag{7}$$

To project the sequence into a fixed dimension and output the representation \mathcal{N}_{sup} of \mathcal{V}_{sup} for graph optimization, a *self-attention* is utilized to summarize the contextual information:

$$j_s = \text{softmax}(\text{MLP}(S_{ca})) \in \mathbb{R}^{l_s \times 1} \tag{8}$$

$$\mathcal{N}_{sup} = j_s^\top S_{ca} \in \mathbb{R}^{1 \times 2d} \tag{9}$$

In addition to the query-aware support documents, the *co-attention* and *self-attention* are used to generate query-aware node representations from other sequential representations.

3.4.2. Message Passing in the Heterogeneous Graph Attention Network

We present messaging passing in the heterogeneous graph attention network for reading within multiple relations in diverse nodes. The input of this module is a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ and node representations $\mathcal{N} = \{n_1, n_2, \dots, n_r\} \in \mathbb{R}^{1 \times 2d}$, where r is the number of nodes. Initially, a shared weight matrix \mathbf{W}_n is applied to \mathcal{N} , then the attention coefficients and nodes attention coefficients are computed as

$$e_{ij} = \text{MLP}(\mathbf{W}_n n_i || \mathbf{W}_n n_j) \tag{10}$$

$$\alpha_{ij} = \text{softmax}_j(e_{ij}) = \frac{\exp(e_{ij})}{\sum_{k \in \mathcal{N}_i} \exp(e_{ik})} \tag{11}$$

where e_{ij} are the attention coefficients indicating the importance of the features of the node n_j to the node n_i , and α_{ij} is normalized across all structure neighbors \mathcal{N}_i of the node n_i .

The attention mechanism is responsible for selectivity with node interdependence, which enables us to show how the nodes take effect during the reasoning.

Considering the 10 different types of edges defined in Section 3.3.2, we model the relational edges basing on the vanilla GAT [37]:

$$n_i^{l+1} = \frac{1}{\mathcal{K}} \parallel_{\kappa=1}^{\mathcal{K}} \sigma \left(\sum_{j \in \mathcal{N}_i} \sum_{r \in \mathcal{R}_{ij}} \frac{1}{|\mathcal{N}_i^r|} \alpha_{r_{ij}}^{k,l} \mathbf{W}_{r_{ij}}^{k,l} n_j^l \right) \tag{12}$$

where $n_i^l \in \mathbb{R}^{1 \times 2d}$ is the hidden state of the node n_i in the l -th layer, all the GAT layers are parameter-shared, κ is the κ -th head following [15,37], \mathcal{R} is the set of all types of edges in \mathcal{E} , and $\alpha_{r_{ij}}^{k,l}$ are normalized attention coefficients computed by the k -th attention mechanism with relation r , which is presented in [37].

Message passing is a key component of our model. To echo the selectivity of *grandmother cells*, we use the attention mechanism to select (i.e., activate or deactivate) key node pairs in our reasoning graph, and we empirically regard this process as the reading reasoning in the graph.

3.4.3. Gating Mechanism

A previous study [19] showed that GNNs are suffer from the smoothing problem when calculated by stacking many layers, and, thus, we overcome this issue by applying question-aware [27] and general gating mechanisms [38] to optimize the procedure.

$$\mathcal{H}_q = \text{BiLSTM}(H_q) \tag{13}$$

$$w_{ij} = \sigma \left(\mathbf{W}_q^\top [n_i^l \parallel \mathcal{H}_{q_j}] \right) \tag{14}$$

$$\alpha_{ij}^{gate} = \frac{\exp(w_{ij})}{\sum_{k=1}^m \exp(w_{ik})} \tag{15}$$

$$q_i^l = \sum_{j=1}^m \alpha_{ij}^{gate} \mathcal{H}_{q_j} \tag{16}$$

$$\beta_i^l = \sigma(\mathbf{W}_s^\top [q_i^l \parallel n_i^l]) \tag{17}$$

$$\tilde{n}_i^l = \beta_i^l \odot \tanh(p_i^j) + (1 - \beta_i^l) \odot n_i^l \tag{18}$$

where H_q is the query representation given by a dedicated *Bi-LSTM* encoder to keep consistency with the dimension of node features \mathcal{N} , j indicates the order of query words, m is the query length, σ is a sigmoid function, and \odot indicates element-wise multiplication. Then the general gating mechanism is introduced as follows:

$$x_i^l = \sigma(\text{MLP}[\tilde{n}_i^l \parallel n_i^l]) \tag{19}$$

$$n_i^{l+1} = x_i^l \odot \tanh(\tilde{n}_i^l) + (1 - x_i^l) \odot n_i^l \tag{20}$$

3.5. Output Layer

After updating the node representation, we use two multilayer perceptrons, MLP_{can} and MLP_{men} , to transform the node features to prediction scores. All the candidate nodes (*grandmother cells*) \mathcal{N}_{can} and mention nodes \mathcal{N}_{men} from \mathcal{G} are employed to output the prediction score distribution a as

$$a = \gamma \times \text{MLP}_{can}(\mathcal{N}_{can}) + (1 - \gamma) \times \max(\text{MLP}_{men}(\mathcal{N}_{men})) \tag{21}$$

where $\max(\cdot)$ takes the maximum mention node score over MLP_{men} , then the two parts are summed with the effect of a harmonic γ as the final prediction score distribution.

4. Experiments

We present the performance of our model on the QANGAROO [13] dataset and evaluate the performance in detail. Then, the ablation study and the visualization will demonstrate the benefit of the model. Finally, a case study shows the relationship between the output of the answer from the models and human reading results.

4.1. Dataset for Experiments

QANGAROO is a multi-hop MRC dataset containing two independent datasets, WIKIHOP and MEDHOP, from the open-domain field and molecular biology field, respectively. Both WIKIHOP and MEDHOP were divided into three subsets: the training set, development set, and undisclosed test set, which is used for official evaluation. The dataset sizes are shown in Table 2.

Table 2. Dataset size of WIKIHOP and MEDHOP.

	Training	Development	Test	Total
WikiHop	43,738	5129	2451	51,318
MedHop	1620	342	546	2508

WIKIHOP was created from WIKIPEDIA (as the document corpus) and WIKIDATA (as structured knowledge triples). A sample from the dataset is shown in Figure 3a. In this sample, the query (*located_in_the_administrative_territorial_entity, hampton_wick_war_memorial, ?*) requires us to answer the administrative territory of the *Hampton Wick War Memorial*. To predict it, a named recognition entity *Hampton Wick* is extracted from the seventh support document, and it links to the same tokens in the zeroth support document where the correct candidate answer appears as well. The reasonable clue chain *Hampton Wick War Memorial* \leftrightarrow *Hampton Wick # 1* \leftrightarrow *Hampton Wick # 2* \leftrightarrow *London Borough of Richmond upon Thames* presents the procedure of our model for the multi-hop MRC task.

To validate whether the dataset can be consistent with the formalization of the multi-hop MRC, the dataset founder asked human annotators to evaluate the samples in the WIKIHOP development and test sets. For each sample in the two sets, at least three annotators participated in the evaluation, and they were required to answer three questions:

- Whether they knew the fact before;
- Whether the fact follows from the texts (with options *follows*, *likely*, and *not follows*);
- Whether multiple documents are required to answer the question.

All the samples in the test set were human-selected and were labeled by the majority of annotators with *follows* and *multiple documents required*. Annotators merely noted the samples in the development set without the selection.

The MEDHOP dataset was constructed using the DRUGBANK as certain knowledge. Then the creators extracted the research paper abstracts from MEDLINE—the online medical literature search and analysis system and the bibliographic database of the National Library of Medicine of the USA—as a corpus, and the aim is to predict the drug–drug interaction (DDI) after reading the texts. The purpose of applying multi-hop methods in this prediction is to find and combine individual observations that can suggest previously unobserved DDI from inferring and reasoning the prior public knowledge in contents rather than some costly experiments. The only query type is *interacts_with*. A sample given in [13] is illustrated in Figure 3b and note that accession numbers replace the medical proper nouns (e.g., DB00007, DB06825, DB00316) rather than the names of drugs and human proteins (e.g., leuprolide, triptorelin, acetaminophen) in practice.

<p>Question: (located_in_the_administrative_territorial_entity, hampton_wick_war_memorial, ?)</p> <p>Document #7: The Hampton Wick War Memorial is situated on the Hampton Wick side of Kingston Bridge, in the south - west suburbs of London, England...</p> <p>Document #0: Hampton Wick, formerly a village, is a Thames-side area of the London Borough of Richmond upon Thames, which ... It is buffered by a public green space, one of the Royal Parks of London from Hampton and Hampton Hill... ...</p> <p>Candidates: {hampton, hampton hill, ..., london borough of richmond upon thames, Scotland}</p> <p>Answer: london borough of richmond upon thames</p>	<p>Question: (Leuprolide, interacts_with, ?)</p> <p>Document #1: Leuprolide ... elicited a long-lasting potentiation of excitatory postsynaptic currents... GnRH receptor-induced synaptic potentiation was blocked ... by Progonadoliberein-1, a specific GnRH receptor antagonist...</p> <p>Document #2: ... our research to study the distribution, co-localization of Urofollitropin and its receptor[,] and co-localization of Urofollitropin and GnRH receptor...</p> <p>Support #3: Analyses of gene expression demonstrated a dynamic response to the Progonadoliberein-1 superagonist Triptorelin. ...</p> <p>Candidates: {Triptorelin, Urofollitropin}</p> <p>Answer: Triptorelin</p>
(a)	(b)

Figure 3. Samples of WIKIHOP and MEDHOP. Subject entities, reasoning entities, mention entities, and candidate entities are shown in gray, orange, green, and blue colors, respectively. The occurrence of the correct answer is shown by a square frame outside. (a) A sample from the WIKIHOP. (b) A sample from the MEDHOP.

4.2. Experiments Settings

We exploited the NLTK [39] toolkit to tokenize the support documents and candidates, then split the query $q = \{s, r, a^*\}$ into relation r and subject entity s . All the named entities matching with candidates C_q were extracted as mention nodes \mathcal{V}_{men} , and the SPACY (<https://spacy.io> (accessed on 21 June 2023)) was used to extract the named entities and noun phrases from texts as reasoning nodes \mathcal{V}_{rea} . We concatenated GloVe [32] and n-gram character embeddings [40] to obtain 400-dimensional word embeddings, which were input to the encoder layer. The out-of-vocabulary words were presented with random vectors. The word embedding was fixed in WIKIHOP experiment and trainable on MEDHOP. We implemented the ClueReader model with PyTorch and PyTorch Geometric [41]. NetworkX [42] was utilized to visualize the reading graph, the weights of node pair weights, and node selections.

4.3. Results and Analyses

In Table 3, we present the performance of ClueReader in the development and test sets of WIKIHOP and MEDHOP and compare it with the performance of published models mainly based on GNNs. Our model improved the accuracy of GCN-based models HDE [28] in the test set from 70.9% to 72.0% and Path-based GCN in the development set from 64.5% to 66.9%, while Path-based GCN using GloVe and ELMo word embeddings surpassed our model by 0.5% in the test set, which confirms that the initial representations of nodes are extremely critical [27]. However, limited by the architecture and computing resources, we did not use powerful contextual word embeddings like ELMo and BERT in our model, which can be further addressed. Compared to the other GNN-based models [24,26,31] and the sequential models [13,43], our model achieved higher accuracy. We are the first to apply the GNN-based model to MEDHOP, although the accuracy was 1.8% lower than BiDAF, we believe that the possible reason was the failure in extracting the reasoning nodes of the SPACY toolkit, which means the bridge entities were incomplete.

To analyze the scalability of our model, we divided the development set into six groups according to the number of support documents and then determined the accuracy in each group. The grouped accuracies on WIKIHOP are shown in Figure 4. *ClueReader* achieved competitive results: 73.59% and 63.57% in the groups of (1–10) and (11–20), with a total of 4039 samples accounting for 95% of the development set. The lowest accuracy of 55.74% was for the group (41–50). However, it increased to 62.5% in the group (51–62), which shows the scalability of our model is effective. The grouped accuracies on MEDHOP are shown in Figure 5, and they are quite competitive. The highest and second-highest accuracies of 60.00% and 51.85% are in (31–40) and (21–30) groups, respectively, and the lowest and second-lowest accuracies of 0% and 35.59% are in (1–10) and (51–62) groups, respectively. In particular, the result in the (51–64) group on MEDHOP is against the group (51–62) on WIKIHOP, which implies that we must concentrate on the difference between the open-domain and molecular textual contexts. The results in the different number of support documents show the contribution of our model to the scalability of the multi-hop MRC tasks.

Table 3. Performance of the proposed *ClueReader* in the development and test sets of WIKIHOP and MEDHOP, and comparisons with other published approaches on the leaderboard.

Single Models	WikiHop Accuracy (%)		MedHop Accuracy (%)	
	Dev	Test	Dev	Test
Coref-GRU [43]	56.0	59.3	-	-
MHQA-GRN [31]	62.8	65.4	-	-
Entity-GCN [24]	64.8	67.6	-	-
HDE [28]	68.1	70.9	-	-
BAG [26]	66.5	69.0	-	-
Path-based GCN [27]	64.5	-	-	-
Document-cue [13]	-	36.7	-	44.9
FastQA [13]	-	25.7	-	23.1
TF-IDF [13]	-	25.6	-	9.0
BiDAF [13]	-	42.9	-	47.8
<i>ClueReader</i>	66.5	72.0	48.2	46.0

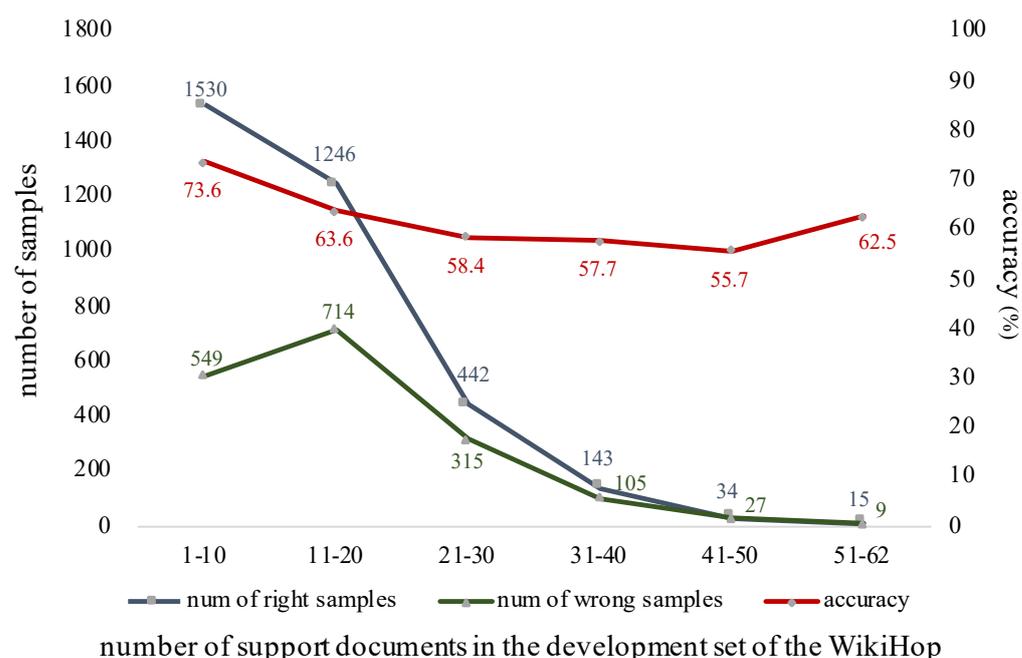


Figure 4. Statistics of the model performance with different numbers of support documents on the WIKIHOP development set.

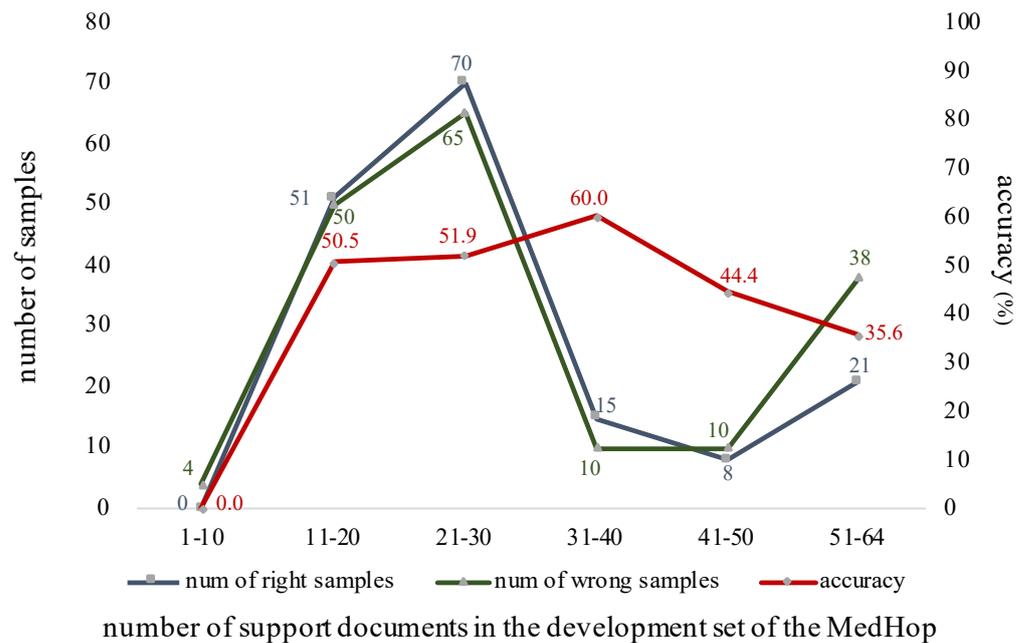


Figure 5. Statistics of the model performance with different numbers of support documents on the MEDHOP development set.

As mentioned above, the WIKIHOP development set had the consistency between facts and documents annotated. To determine whether multiple documents are required to reason the question, we split our models into five categories as follows. In each category, all three annotators annotated: (1) *requires multiple documents and follows fact*; (2) *requires single document and follows fact*; (3) *requires multiple documents and likely follows fact*; (4) *requires single document and likely follows fact*; (5) *not follows is not given*. The performance of our model is presented in Table 4. We observe that *ClueReader* had the best performance of 74.9% in the samples which follow the facts and require multiple passages. This phenomenon proves the effectiveness of the model in pure multi-hop MRC tasks. It achieved the second-best result of 74.0% in samples following the facts and requiring a single document, which supports that *ClueReader* is also effective in single-passage MRC tasks. Further, we believe that authenticity can seriously impact the accuracy of our prediction. The categories associated with *may not follow the fact* achieved the worse results, of 71.4%, 71.4%, and 71.5%, respectively, in the groups of *likely follows the fact (single document and multiple documents)* and *“not follows” is not given*. The same analysis is infeasible in the development set of MEDHOP since the document complexity and the number of documents per sample are significantly larger.

Table 4. Performance on the WIKIHOP development set.

Annotation		Accuracy (%)
follows fact	<i>requires multiple documents</i>	74.9
	<i>requires single document</i>	74.0
likely follows fact	<i>requires multiple documents</i>	71.4
	<i>requires single document</i>	71.4
<i>not follows is not given</i>		71.5

4.4. Ablation Study

We proposed five types of nodes in \mathcal{G} , and to analyze how they reasoned, we removed the edges with specific connections and isolated the nodes to evaluate the performance in the subset of the WIKIHOP development set; that is, *not follows* was not annotated. Moreover, we tested the model without the message passing in \mathcal{G} . The ablated performance is shown in Table 5.

Table 5. Ablation performance on the QANGAROO development set.

Model	Accuracy (%)			
	WIKIHOP	Δ	MEDHOP	Δ
Full Model	71.45	-	48.25	-
<i>w/o GAT</i>	52.69	18.76	37.72	10.53
<i>w/o N_{sub}</i>	70.95	0.5	47.37	0.88
<i>w/o N_{men}</i>	63.34	8.11	4.97	43.28
<i>w/o N_{rea}</i>	70.77	0.68	47.37	0.88
<i>w/o N_{sup}</i>	62.02	9.43	48.54	-0.29
<i>w/o N_{can}</i>	65.87	5.58	44.77	3.48

On WIKIHOP, the proposed heterogeneous graph attention network was the most effective component of *ClueReader*. Without its contribution, the accuracy decreased by 18.76%. After blocking the nodes by groups, we observed that the support nodes contributed 9.43% absolutely, the mention nodes dedicated 8.11%, and the candidate nodes contributed 5.58%. Regarding the reasoning and subject nodes, we considered the small quantities contained in the graph leading to low status in contributions. However, we observed considerably different performances between WIKIHOP and MEDHOP. As the results show in Table 5, the most effective part of the model is mention nodes. When we blocked the mention nodes in the graph, the accuracy decreased significantly, by 43.28%, and the graph reasoning contributed 10.53% to accuracy. Meanwhile, support nodes had negative effects on the prediction, a decrease of 0.29%, which is diametrically opposite the performance on the WIKIHOP development subset.

In Table 6, we present the model performances with different hyperparameters, especially the number of stacked GAT layers (the number of hops) and the weight of *grandmother cells*. The number of GAT layers controls how many parameter-sharing GAT layers should be involved in the reasoning graph. On WIKIHOP, we obtained the highest accuracy (66.5%) when we stacked the graph with five layers, and the model with three or four GAT layers had poorer performance (57.8% or 58.5%, respectively). With six GAT layers, the accuracy dropped 2.3% compared with the best performance.

Table 6. Ablation studies of hyperparameters of GAT layers and weights of *grandmother cells* in reasoning graph predictions.

Hyperparameters	Value	Acc. of WIKIHOP	Acc. of MEDHOP
l	3	57.8	42.4
	4	58.5	43.3
	5	66.5	48.2
	6	64.2	45.0
γ	0	59.7	42.7
	0.5	66.1	44.2
	1.0	66.5	48.2
	1.5	59.1	43.3

Furthermore, as the final prediction illustrated in Equation (21), γ coordinates the mention nodes and the candidate nodes *grandmother cells*; we present the model performances with different γ settings in Table 6. The best performance was with γ set to 1. However, if we gave it too much weight, that is, $\gamma = 1.5$, the accuracy decreased by 7.4%, which is even worse than when we set γ to 0 (59.7%), which convinces us that we should not ignore the effect of much larger networks behind *grandmother cells*. We observed similar phenomena with different hyperparameter settings on MEDHOP. When the number of hops was five, and γ was 1, the model performed best at approximately 48.2%. We suspect that when a few GAT layers are stacked, the messages of nodes cannot pass sufficiently among the reasoning graph. When too many GAT layers are stacked, the graph over-smoothing problem leads to a drop in accuracy. We also empirically observed that models with higher γ may lose semantic information from context resulting in reduced prediction accuracy, which also fits the concept of *grandmother cells* in that before the final predicting determination, a huge background network calculation should be performed implicitly.

4.5. Visualization

Compared to spectral GNN-based reading approaches, our proposed heterogeneous reasoning graph *ClueReader* is a non-spectral approach, which allows us to analyze how the nodes interact with each other in various relations and how the connections take effect between nodes. We visualize the predictions in our heterogeneous reasoning graph on WIKIHOP and MEDHOP in Figures 6 and 7, respectively. Different types of nodes are shown in different colors (subject nodes are gray, reasoning nodes are orange, mention nodes are green, candidate nodes are blue, and support nodes are red), and their edges, which reflect selections of node pairs, are shown in different thickness lines. The thicker the edges, the more important they learn from training. Considering that the answer determination should not only be inferred by the weight edges but also from the output layer projected from the representations of the nodes to $\mathbb{R}^{1 \times 2d}$ and accumulated score from \mathcal{N}_{can} and \mathcal{N}_{men} , we use the transparency of the nodes to respond to the outputs: the darker the nodes, the higher the values output from the output layer. Owing to the output values being quite different, some mention and candidate nodes are almost transparent. The weight graph provides the evidence during reading and the analysis of DDI. It passes the messages according to the concept of *grandmother cells* that not only one node becomes effective, but the cluster behind it plays a synergistic effect. We learn more about our model through visualization. For instance, the node transparency differentiation on MEDHOP is significantly lower than WIKIHOP, which indicates that the drug features are not sufficiently learned, leading to the convergence of node features and increased classification prediction difficulty. This issue can be further addressed.

To better understand the model predictions and contribute to further study, we generate HTML files of samples as shown in Figure 8 and analyze whether the named entities contained in the max-score nodes can make sense from the perspective of human answering after reading. Please refer to our project website (<https://github.com/cluereader/cluereader.github.io> (accessed on 21 June 2023)) for more visualization samples in HTML files.

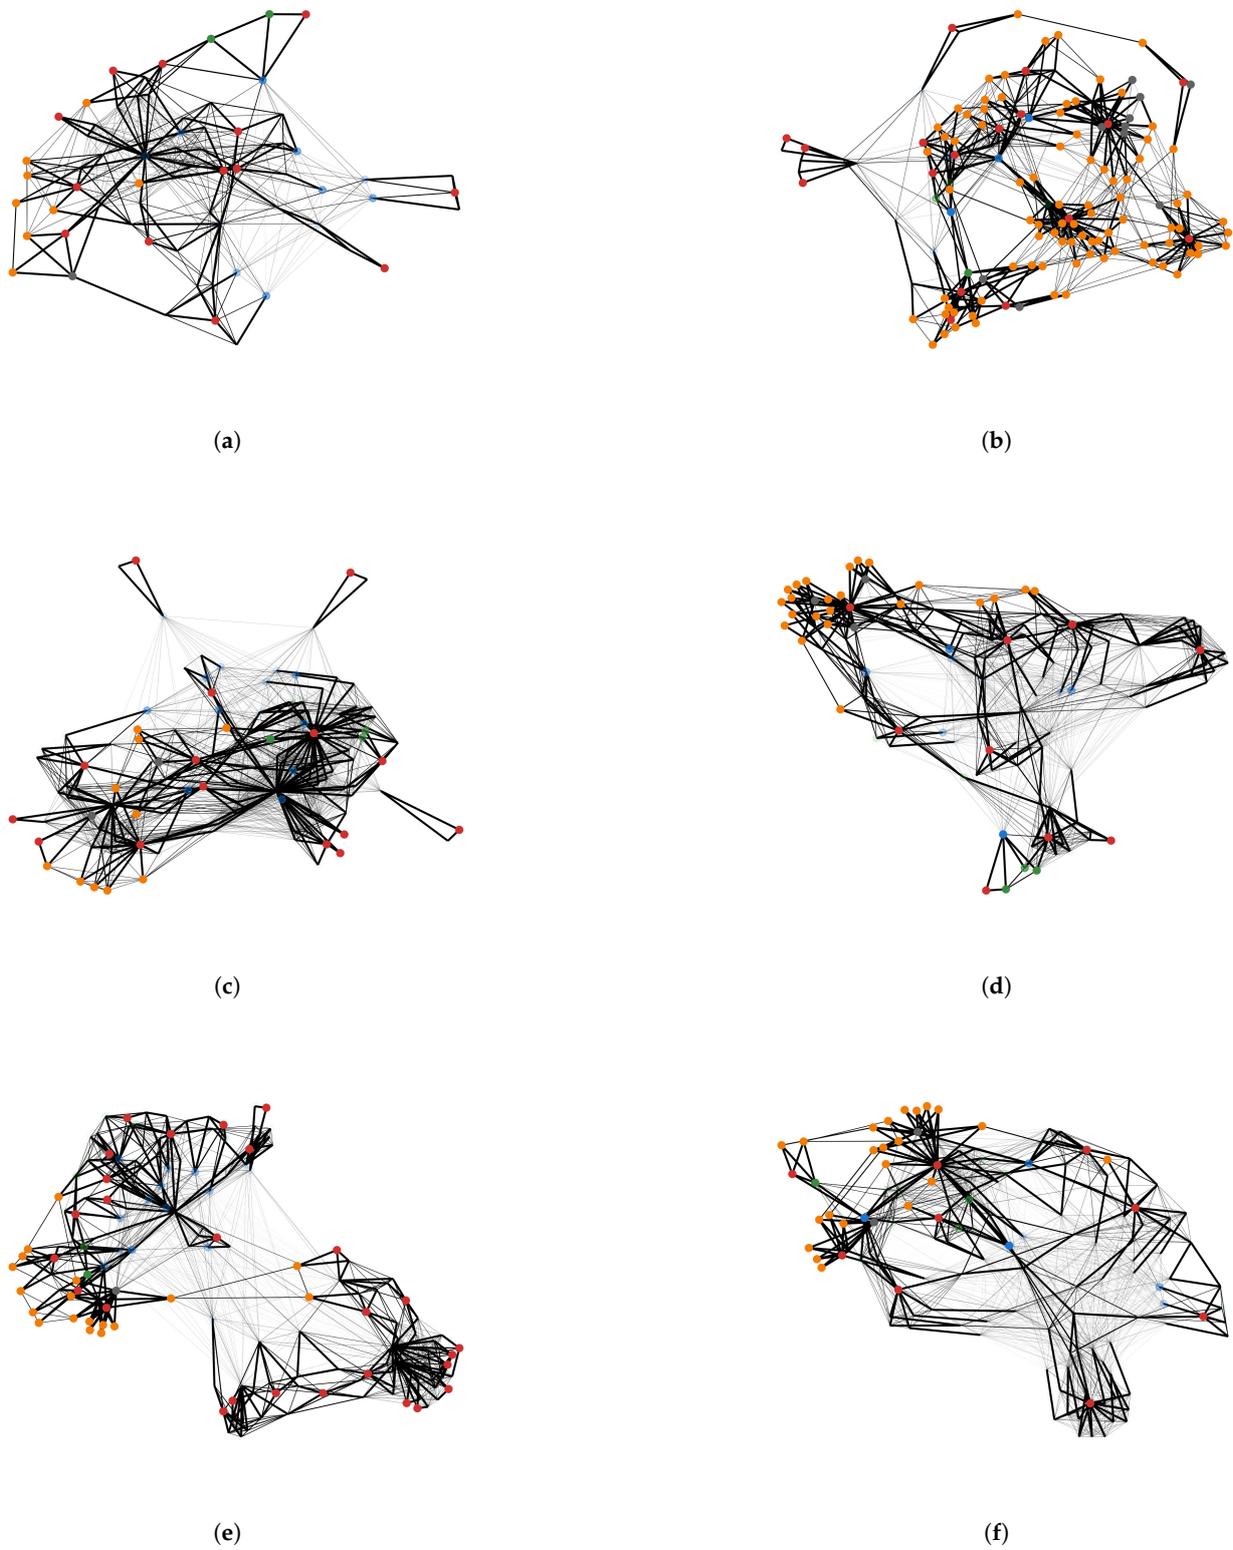


Figure 6. Visualizations of reasoning graphs on the WIKIHOP development set that are correctly answered. A thicker edge corresponds to a higher attention weight, and darker green nodes or darker blue nodes represent higher output values among the same type of nodes. (a–f) Visualized samples from the WIKIHOP development set.

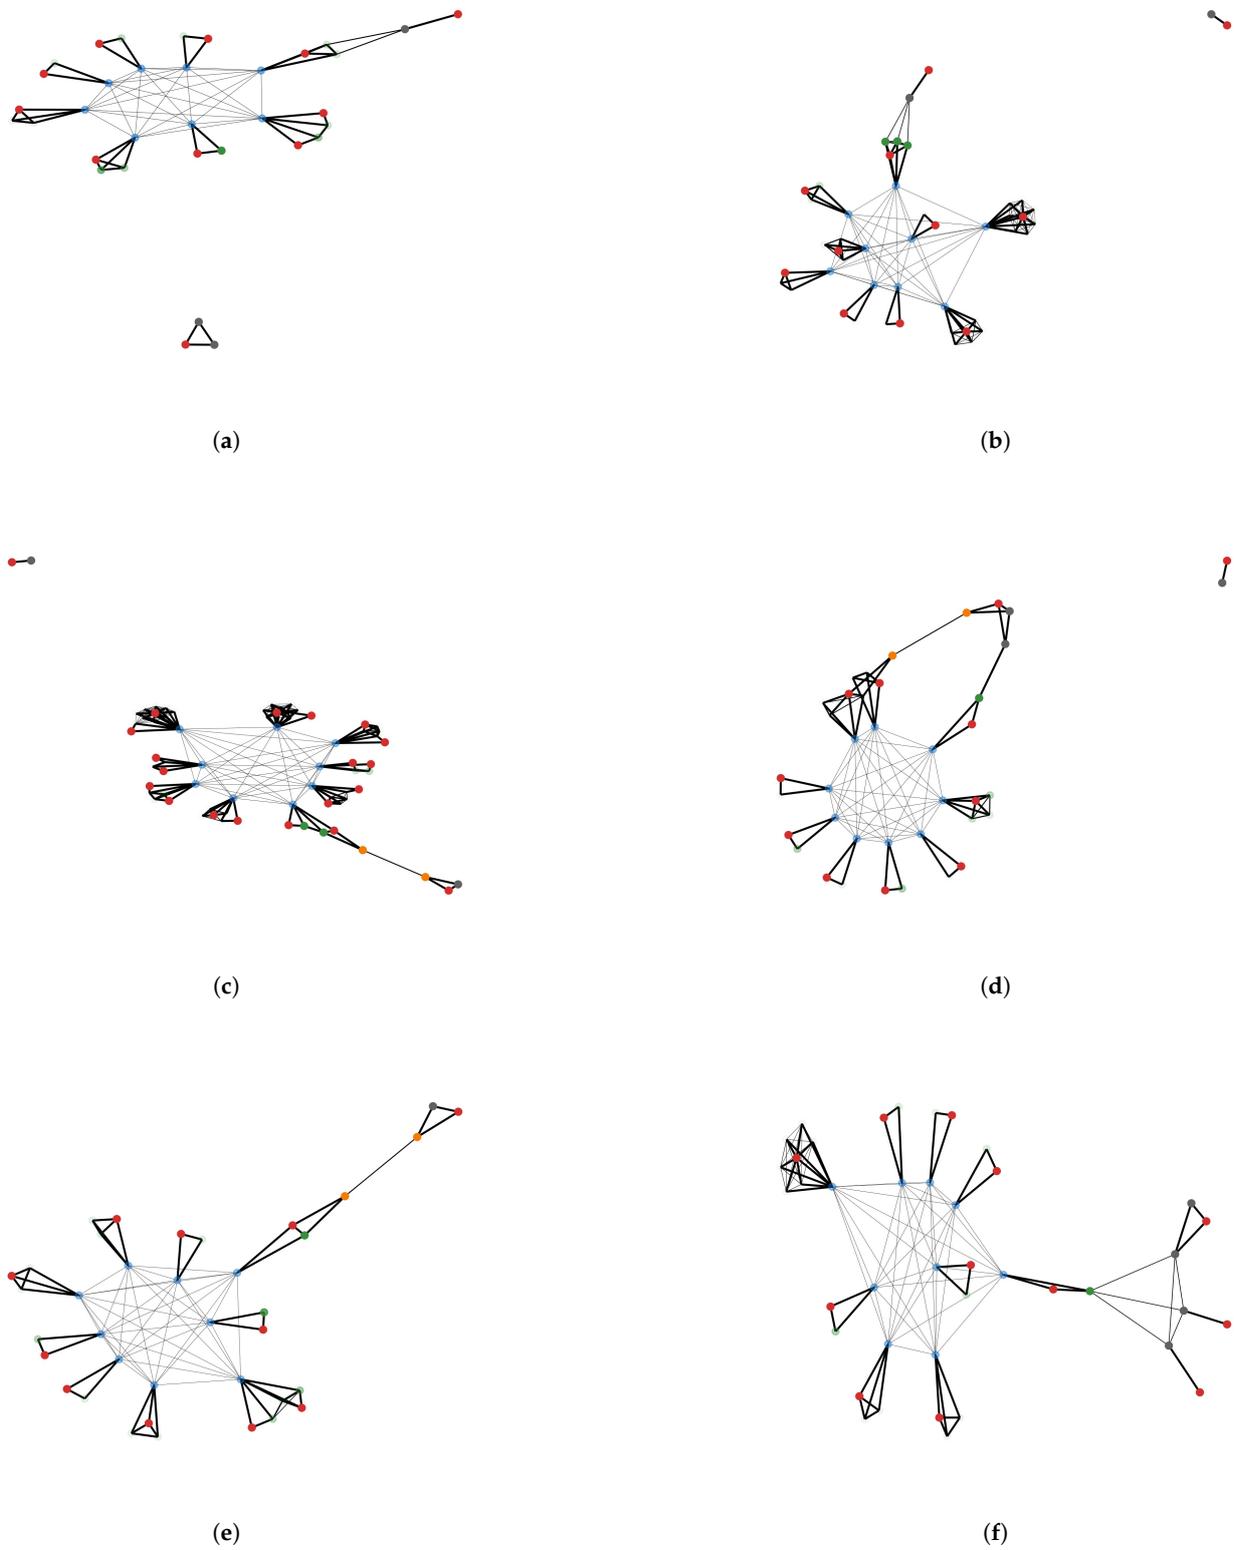


Figure 7. Visualizations of reasoning graphs on the MEDHOP development set that are correctly answered. A thicker edge corresponds to a higher attention weight, and darker green nodes or darker blue nodes represent higher output values among the same type of nodes. (a–f) Visualized samples from the MEDHOP development set.

ID WH_dev_543
Query located_in_the_administrative_territorial_entity queensville
Candidates alberta, alpine, calgary, canada, capital region, etc.,
Answer regional municipality of york

Documents

0

FOREST MEN (short for "Freedom MEN Organisation for the Right to Enjoy Smoking Tobacco") is a United Kingdom MEN political pressure group which campaigns against tobacco control activity.

1

East MEN East Gwillimbury REA is a town on the East MEN Holland MEN River MEN in the Regional Municipality of York MENMAX Regional Municipality of York MEN of MEN York. It is part of MEN the Greater Toronto MEN Area of MEN southern Ontario MEN Ontario REA, in Canada MEN Canada REA. It was formed by the amalgamation of MEN the Township of MEN East MEN East Gwillimbury REA with all the previously incorporated villages and hamlets within the township. The main centres REA in East MEN East Gwillimbury REA are the villages REA of MEN Holland MEN Holland Landing REA, Queensville SUB, Sharon MEN Sharon REA, and Mount Albert REA. The Civic Centre (municipal offices) are located along Leslie Street in Sharon MEN Sharon REA. The northernmost interchange of MEN Highway 404 is at the North MEN edge of MEN East MEN East Gwillimbury REA, just south MEN of MEN Ravenshoe Road MEN. The hamlets of MEN Holt and Brown Hill MEN are also within town limits.

Figure 8. Generated HTML file of sample # 543 in WIKIHOP development set. The mark **MENMAX** means the final output of MLP_{men} . For more details, please refer to https://clureader.github.io/WH_dev_543.html (accessed on 21 June 2023).

5. Conclusions

We present *ClueReader*, a heterogeneous graph attention network for multi-hop MRC, which is inspired by the concept of *grandmother cells* from cognitive neuroscience. The network contains several clue-reading paths from the subject of the question and ends with candidate entities. We use reasoning and mention nodes to complete the process and use support nodes to add supernumerary semantic information. We apply our methodology on QANGAROO, a multi-hop MRC dataset, and the official evaluation supports the effectiveness of our model in open-domain QA and the molecular biology domain. Several potential issues could be further addressed, such as introducing intermediate supervision signals during the semi-supervised graph learning, the enhancement of using external knowledge, and dedicated word-embedding methodology in the medical context, which are possible to improve the model performance in multi-hop MRC tasks.

Author Contributions: P.G.: conceptualization, methodology, investigation, writing—review & editing, and software. F.G.: writing—original draft, writing—review & editing, visualization, and software. P.W.: conception, design, and resources. J.-C.N.: software and data curation. F.W.: supervision and validation. H.F.: formal analysis and validation. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded in part by the China Postdoctoral Science Foundation under Grant 2023M732022, in part by the Excellent Young Scientists Fund Program of Qufu Normal University under Grant 167/602801, in part by the Shandong Provincial Natural Science Foundation under Grant ZR2021QF061 and ZR2022MF353 and in part by the Guangdong Provincial Basic and Applied Basic Research Foundation under Grant 2020A1515010706.

Data Availability Statement: The data presented in this study are available upon request from the corresponding author.

Acknowledgments: The authors would like to thank the UCL machine reading group that created the QANGAROO dataset and their help in evaluating *ClueReader*.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Wang, Y.; Liu, K.; Liu, J.; He, W.; Lyu, Y.; Wu, H.; Li, S.; Wang, H. 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), Melbourne, Australia, 15–20 July 2018; pp. 1918–1927.
2. Dai, Y.; Fu, Y.; Yang, L. A multiple-choice machine reading comprehension model with multi-granularity semantic reasoning. *Appl. Sci.* **2021**, *11*, 7945. [[CrossRef](#)]
3. Hassabis, D.; Kumaran, D.; Summerfield, C.; Botvinick, M. Neuroscience-inspired artificial intelligence. *Neuron* **2017**, *95*, 245–258. [[CrossRef](#)]
4. Page, M. Connectionist modelling in psychology: A localist manifesto. *Behav. Brain Sci.* **2000**, *23*, 443–467. [[CrossRef](#)] [[PubMed](#)]
5. Dehaene, S. *Reading in the Brain: The New Science of How We Read*; Penguin: London, UK, 2010.
6. Quiroga, R.Q.; Reddy, L.; Kreiman, G.; Koch, C.; Fried, I. Invariant visual representation by single neurons in the human brain. *Nature* **2005**, *435*, 1102–1107. [[CrossRef](#)]
7. Battaglia, P.W.; Hamrick, J.B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; et al. Relational inductive biases, deep learning, and graph networks. *arXiv* **2018**, arXiv:1806.01261.
8. Yu, W.; Chang, T.; Guo, X.; Wang, M.; Wang, X. An interaction-modeling mechanism for context-dependent Text-to-SQL translation based on heterogeneous graph aggregation. *Neural Netw.* **2021**, *142*, 573–582. [[CrossRef](#)] [[PubMed](#)]
9. Seo, M.; Kembhavi, A.; Farhadi, A.; Hajishirzi, H. Bidirectional Attention Flow for Machine Comprehension. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
10. Cui, Y.; Chen, Z.; Wei, S.; Wang, S.; Liu, T.; Hu, G. Attention-over-Attention Neural Networks for Reading Comprehension. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Vancouver, BC, Canada, 30 July–4 August 2017; pp. 593–602.
11. Zhao, Y.; Wang, L.; Wang, C.; Du, H.; Wei, S.; Feng, H.; Yu, Z.; Li, Q. Multi-granularity heterogeneous graph attention networks for extractive document summarization. *Neural Netw.* **2022**, *155*, 340–347. [[CrossRef](#)] [[PubMed](#)]
12. Li, S.; Sun, C.; Liu, B.; Liu, Y.; Ji, Z. Modeling Extractive Question Answering Using Encoder-Decoder Models with Constrained Decoding and Evaluation-Based Reinforcement Learning. *Mathematics* **2023**, *11*, 1624. [[CrossRef](#)]
13. Welbl, J.; Stenetorp, P.; Riedel, S. Constructing datasets for multi-hop reading comprehension across documents. *Trans. Assoc. Comput. Linguist.* **2018**, *6*, 287–302. [[CrossRef](#)]
14. Kenton, J.D.M.W.C.; Toutanova, L.K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the NAACL-HLT, Minneapolis, MN, USA, 2–7 June 2019; pp. 4171–4186.
15. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 1–11.
16. Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv* **2019**, arXiv:1907.11692.
17. Beltagy, I.; Peters, M.E.; Cohan, A. Longformer: The long-document transformer. *arXiv* **2020**, arXiv:2004.05150.
18. Razeghi, Y.; Logan, R.L., IV; Gardner, M.; Singh, S. Impact of pretraining term frequencies on few-shot reasoning. *arXiv* **2022**, arXiv:2202.07206.
19. Kipf, T.N.; Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
20. Zhang, C.; Zha, D.; Wang, L.; Mu, N.; Yang, C.; Wang, B.; Xu, F. Graph Convolution Network over Dependency Structure Improve Knowledge Base Question Answering. *Electronics* **2023**, *12*, 2675. [[CrossRef](#)]
21. Ding, M.; Zhou, C.; Chen, Q.; Yang, H.; Tang, J. Cognitive Graph for Multi-Hop Reading Comprehension at Scale. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; pp. 2694–2703.
22. Evans, J.S.B. Heuristic and analytic processes in reasoning. *Br. J. Psychol.* **1984**, *75*, 451–468. [[CrossRef](#)]
23. Sloman, S.A. The empirical case for two systems of reasoning. *Psychol. Bull.* **1996**, *119*, 3. [[CrossRef](#)]
24. De Cao, N.; Aziz, W.; Titov, I. Question Answering by Reasoning Across Documents with Graph Convolutional Networks. In Proceedings of the 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics. Association for Computational Linguistics, Minneapolis, MN, USA, 3–7 June 2019; pp. 2306–2317.
25. Peters, M.E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; Zettlemoyer, L. Deep Contextualized Word Representations. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). Association for Computational Linguistics, New Orleans, LA, USA, 1–6 June 2018; pp. 2227–2237.
26. Cao, Y.; Fang, M.; Tao, D. BAG: Bi-directional Attention Entity Graph Convolutional Network for Multi-hop Reasoning Question Answering. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Minneapolis, MN, USA, 3–7 June 2019; pp. 357–362.
27. Tang, Z.; Shen, Y.; Ma, X.; Xu, W.; Yu, J.; Lu, W. Multi-hop reading comprehension across documents with path-based graph convolutional network. In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence, Yokohama, Japan, 7–15 January 2021; pp. 3905–3911.

28. Tu, M.; Wang, G.; Huang, J.; Tang, Y.; He, X.; Zhou, B. Multi-hop Reading Comprehension across Multiple Documents by Reasoning over Heterogeneous Graphs. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; pp. 2704–2713.
29. Jia, M.; Liao, L.; Wang, W.; Li, F.; Chen, Z.; Li, J.; Huang, H. Keywords-aware dynamic graph neural network for multi-hop reading comprehension. *Neurocomputing* **2022**, *501*, 25–40. [[CrossRef](#)]
30. Zhang, Y.; Meng, F.; Zhang, J.; Chen, Y.; Xu, J.; Zhou, J. MKGN: A Multi-Dimensional Knowledge Enhanced Graph Network for Multi-Hop Question and Answering. *IEICE Trans. Inf. Syst.* **2022**, *105*, 807–819. [[CrossRef](#)]
31. Song, L.; Wang, Z.; Yu, M.; Zhang, Y.; Florian, R.; Gildea, D. Evidence integration for multi-hop reading comprehension with graph neural networks. *IEEE Trans. Knowl. Data Eng.* **2020**, *34*, 631–639. [[CrossRef](#)]
32. Pennington, J.; Socher, R.; Manning, C.D. Glove: Global vectors for word representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 25–29 October 2014; pp. 1532–1543.
33. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)]
34. Zhou, P.; Shi, W.; Tian, J.; Qi, Z.; Li, B.; Hao, H.; Xu, B. Attention-based bidirectional long short-term memory networks for relation classification. In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Berlin, Germany, 7–12 August 2016; pp. 207–212.
35. Tang, H.; Li, H.; Liu, J.; Hong, Y.; Wu, H.; Wang, H. DuReader_robust: A Chinese Dataset Towards Evaluating Robustness and Generalization of Machine Reading Comprehension in Real-World Applications. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), Virtual, 1–6 August 2021; pp. 955–963.
36. Zhong, V.; Xiong, C.; Keskar, N.S.; Socher, R. Coarse-grain Fine-grain Coattention Network for Multi-evidence Question Answering. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
37. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; Bengio, Y. Graph Attention Networks. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April– 3 May 2018.
38. Gilmer, J.; Schoenholz, S.S.; Riley, P.F.; Vinyals, O.; Dahl, G.E. Neural message passing for quantum chemistry. In Proceedings of the International Conference on Machine Learning, PMLR, Sydney, Australia, 6–11 August 2017; pp. 1263–1272.
39. Bird, S. NLTK: The natural language toolkit. In Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions, Sydney, Australia, 17–18 July 2006; pp. 69–72.
40. Hashimoto, K.; Xiong, C.; Tsuruoka, Y.; Socher, R. A Joint Many-Task Model: Growing a Neural Network for Multiple NLP Tasks. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Copenhagen, Denmark, 9–11 September 2017; pp. 1923–1933.
41. Fey, M.; Lenssen, J.E. Fast graph representation learning with PyTorch Geometric. *arXiv* **2019**, arXiv:1903.02428.
42. Hagberg, A.; Swart, P.; Chult, D.S. *Exploring Network Structure, Dynamics, and Function Using NetworkX*; Technical Report; Los Alamos National Laboratory: Los Alamos, NM, USA, 2008.
43. Dhingra, B.; Jin, Q.; Yang, Z.; Cohen, W.; Salakhutdinov, R. Neural Models for Reasoning over Multiple Mentions Using Coreference. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers), New Orleans, LA, USA, 1–6 June 2018; pp. 42–48.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.