

# Interactive Machine Comprehension with Dynamic Knowledge Graphs

Xingdi Yuan

Microsoft Research, Montréal  
eric.yuan@microsoft.com

## Abstract

Interactive machine reading comprehension (iMRC) is machine comprehension tasks where knowledge sources are partially observable. An agent must interact with an environment sequentially to gather necessary knowledge in order to answer a question. We hypothesize that graph representations are good inductive biases, which can serve as an agent’s memory mechanism in iMRC tasks. We explore four different categories of graphs that can capture text information at various levels. We describe methods that dynamically build and update these graphs during information gathering, as well as neural models to encode graph representations in RL agents. Extensive experiments on iSQuAD suggest that graph representations can result in significant performance improvements for RL agents.<sup>1</sup>

## 1 Introduction

Machine reading comprehension (MRC) has gathered wide interest from the NLP community in recent years. It serves as a way to benchmark a system’s ability to understand and reason over natural language. Typically, given a knowledge source such as a document, a model is required to read through the knowledge source to answer a question about some information contained therein. In the extractive QA paradigm, in particular, answers are typically sub-strings of the knowledge source (Rajpurkar et al., 2016; Trischler et al., 2016a; Yang et al., 2018a). Models are thus required to select a span from the knowledge source as their prediction.

A recent line of work known as interactive machine reading comprehension (iMRC) features interactive language learning and knowledge acquisition (Yuan et al., 2020; Ferguson et al., 2020). It shifts the focus of MRC research towards a more realistic setting where the knowledge sources (environments) are *partially observable*. Under this

setting, agents must iteratively interact with the environment to discover necessary information in order to answer the questions. The sequence of interactions between an agent and the environment may resemble the agent’s reasoning path, rendering a higher level of interpretability in the agent’s behaviour. The trajectories of interactions can also be seen as procedural knowledge, which potentially brings agents extra generalizability (humans do not necessarily know the answer to a question immediately, but they know the procedure to search it). Compared to many static MRC datasets, where the entire knowledge source (e.g., a paragraph in SQuAD) is presented to the model immediately, the iMRC setting may alleviate the risk of learning shallow pattern matching (Sugawara et al., 2018; Sen and Saffari, 2020).

On a parallel track, there have been a plethora of studies that leverage graphs in MRC. Multiple linguistic features have been explored to help construct graphs, such as coreference (Dhingra et al., 2018; Song et al., 2018), entities (De Cao et al., 2019; Qiu et al., 2019; Tu et al., 2020) and semantic roles (Zheng and Kordjamshidi, 2020). In related areas such as vision- and text-based games, prior works also attempt to build implicit and explicit graphs to encode data from various types of modalities (Johnson, 2017; Ammanabrolu and Riedl, 2019; Kipf et al., 2020; Adhikari et al., 2020). All of these works, covering domains from static MRC to sequential decision making, suggest that graph representations can facilitate model learning. This gives us a strong motivation to leverage graph representations in the iMRC setting.

We hypothesize that graph representations are good inductive biases, since they can serve naturally as a memory mechanism to help RL agents tackle partial observability in iMRC tasks. We develop an agent that can dynamically update new information into a graph representation at every step, the agent integrates information from both

<sup>1</sup>We release code and data at [https://github.com/xingdi-eric-yuan/imrc\\_graph\\_public](https://github.com/xingdi-eric-yuan/imrc_graph_public)

text and graph modalities to make decisions. The main contributions of this work are as follows:

1. We propose four categories of graph representations, each capturing text information from a unique perspective; we demonstrate how to generate and maintain the graphs dynamically in the iMRC tasks.
2. We extend the RL agent proposed in (Yuan et al., 2020) by adding a graph encoding mechanism and a recurrent memory mechanism.
3. We conduct extensive experiments and show that the proposed graph representations can greatly boost agent’s performance on iMRC tasks.

## 2 Problem Setting

We follow the iMRC setting (Yuan et al., 2020), where given an environment consisting of a *partially observable* document and a question, an agent needs to sequentially interact with the environment to discover necessary information and then answer the question. The iMRC paradigm reformulates existing MRC datasets (e.g., SQuAD) into interactive environments by occluding most parts of their documents. A set of commands are defined to help agents reveal glimpses of the hidden documents.

iMRC can be seen as a controllable simulation to a family of complex real world environments, where knowledge sources are partially observable yet easily accessible by design through interactions. One such example is the Internet, where humans can efficiently navigate through keywords and links to retrieve only the necessary information, rather than reading through the entire collection of websites. While iMRC shares some common properties with multi-step retrieval (Yang et al., 2018a; Zhao et al., 2021) and open-domain QA (Lewis et al., 2020), the focus here is to push the boundaries of information-seeking agents (Bachman et al., 2016) from an RL/navigation perspective.

Formally, an iMRC data-point (game) is a discrete-time, partially observable Markov decision process (POMDP) defined by  $(S, T, A, \Omega, O, R, \gamma)$ . At game step  $t$ , the environment state  $s_t \in S$  represents the semantics and information contained in the full document, as well as which subset of the sentences has been revealed to the agent. The agent perceives text information as its observation,  $o_t \in \Omega$ , which

depends on the environment state with probability  $O(o_t|s_t)$ . The agent issues an action  $a_t \in A$ , resulting in a state transition  $s_{t+1}$  with probability  $T(s_{t+1}|s_t, a_t)$  in the environment (i.e., a new sentence is shown to the agent). Based on its actions, the agent receives rewards  $r_t = R(s_t, a_t)$ . The agent’s objective is to maximize the expected discounted sum of rewards  $E[\sum_t \gamma^t r_t]$ , where  $\gamma \in [0, 1]$  is the discount factor.

**Difficulty Levels** Given a question, only the first sentence of a document is initially exposed to an agent. During information gathering phase, the agent uses the following commands to interact with the environment: 1) **previous** and 2) **next** will jump to the previous or next sentence, respectively; 3) **Ctrl+F QUERY**: jumps to the sentence with the next occurrence of **QUERY**; 4) **stop** terminates the interaction. Whenever the agent issues the **stop** action, or it has exhausted its interaction budget<sup>2</sup>, the information gathering phase is terminated, and the agent needs to answer the question immediately. Thanks to the extractive nature of MRC datasets such as SQuAD, agents can label a span from its observation  $o_t$  as prediction. Note that in order to correctly answer the question, an agent needs to effectively gather necessary information so that its observation  $o_t$  contains the answer as a sub-string.

Yuan et al. (2020) define easy and hard as two difficulty levels. In the easy mode, all four commands are available during information gathering phase; whereas in the hard mode, only **Ctrl+F** and **stop** can be used. Intuitively, in the easy mode, an agent can rely on the **next** command to traverse the entire document, which essentially reduces the problem to learning to stop at the right sentence. In contrast, in the hard mode, the agent is forced to **Ctrl+F** in a smart manner to navigate to potentially informative sentences.

**QUERY Types** Three finer-grained settings are further defined according to the action space of the **Ctrl+F QUERY** command. Specifically, ranging from easy to hard, **QUERY** can be a token extracted from the question ( $q$ ), the concatenation of the question and currently observable sentence ( $q + o_t$ ), or any token selected from the dataset’s vocabulary (*vocab*). From an RL perspective, the sizes of the three settings’ action space can differ by several orders of magnitude (e.g., a question with 10 tokens in  $q$  versus a vocabulary size of 20K in *vocab*).

<sup>2</sup>An agent can at most interact 20 steps.

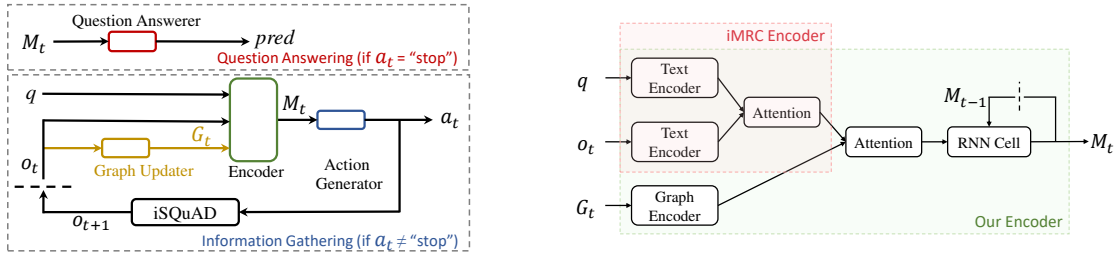


Figure 1: **Left:** an overview of our agent. We propose to use graph representations as an additional input modality to the iMRC agent (Yuan et al., 2020). **Right:** a zoomed in view of our encoder module, extended from iMRC.

### 3 Methodology

In this work, we adopt the agent proposed in (Yuan et al., 2020) as baseline, as shown in Figure 1. We propose to add a novel graph updater module and a graph encoding layer into the pipeline. Specifically, at game step  $t$ , the graph updater takes the text observation  $o_t$  and the graph  $G_{t-1}$  from previous step as input and generates a new graph  $G_t$ . Subsequently, the graph is encoded into hidden states, which is later aggregated with text representations. Note that distinct from fully observed Knowledge Graphs (KGs) in static MRC works, our graphs are dynamically generated, i.e., at every interaction step, our agent can update information from the new observation into its graph representations.

In this section, we will first introduce the key methods we use to generate and update the graph representations. Later on, we will describe a Graph Neural Network (GNN)-based graph encoder which encodes the information carried by the graphs. For the common components shared with iMRC, we refer readers to (Yuan et al., 2020) or Appendix A for detailed information.

**Notations** We denote a graph generated at game step  $t$  as  $G_t = (\mathcal{V}_t, \mathcal{E}_t)$ , where  $\mathcal{V}_t$  and  $\mathcal{E}_t$  represent the set of vertices (nodes) and edges (relations). All graphs are directed by default. For two nodes  $i \in \mathcal{V}$  and  $j \in \mathcal{V}$ , we denote the connection from  $i$  to  $j$  as  $e_{i \rightarrow j} \in \mathcal{E}$ . We represent graph  $G$  as an adjacency tensor, with the size of  $\mathcal{R} \times \mathcal{N} \times \mathcal{N}$ , where  $\mathcal{R}$  and  $\mathcal{N}$  denote the number of relations and nodes, respectively. This tensor can either be binary or real-valued depending on graph type.

#### 3.1 Generating and Updating Graphs

We propose four different graph representations. The four graph types capture distinct aspects of information in text, from lower level linguistic features to high level semantics.

##### 3.1.1 Word Co-occurrence (Rule-based)

$$G \in \{0, 1\}^{\mathcal{R} \times \mathcal{N} \times \mathcal{N}}, \mathcal{V}: \text{words}, \mathcal{E}: \text{sentences.}$$

In word co-occurrence graphs, we connect tokens according to their co-occurrences. We assume that words appear in the same sentence tend to be relevant. Common words across sentences further enable to build more complex graphs where each word is connected with multiple related concepts.

Omitting the notation of game step  $t$  for simplicity, if two tokens  $i \in \mathcal{V}$  and  $j \in \mathcal{V}$  co-occur in a sentence  $s$ , the edge  $e_{i \rightarrow j}$  between  $i$  and  $j$  is defined as  $s$ . Computationally (e.g., for GNNs), the representations of  $i$  and  $j$  are their word embeddings, the representation of the relation  $e_{i \rightarrow j}$  is the sentence encoding of  $s$ . Note in this setting, graphs are symmetrical (i.e.,  $e_{i \rightarrow j} = e_{j \rightarrow i}$ ). Typically, tokens  $i$  and  $j$  can co-occur in multiple sentences. We thus allow multiple connections to appear between two graph nodes, each connection represents a particular sentence where they co-occur.

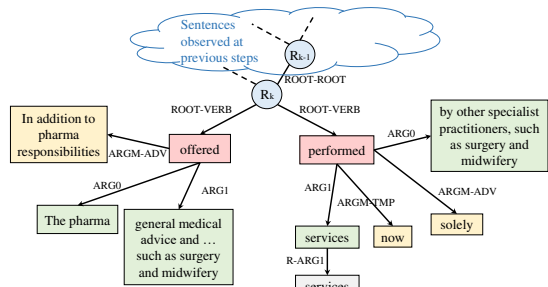
##### 3.1.2 Relative Position (Rule-based)

$$G \in \{0, 1\}^{\mathcal{R} \times \mathcal{N} \times \mathcal{N}}, \mathcal{V}: \text{words}, \\ \mathcal{E}: \text{relative position between words.}$$

In relative position graphs, we aim to capture and embed the word ordering and distance information. This can be seen as capturing a loose form of the *Subject-Verb-Object (SVO)* structure within sentences, without the need of parsing them. The intuition of capturing token position information is supported by the idea of position embeddings in training large-scale language models such as BERT (Devlin et al., 2018; Wang et al., 2021).

In our setting, we first define a window size  $l \in \mathbb{Z}^+$ . For any two tokens  $i$  and  $j$  within a sentence  $s$ , their relation  $e_{i \rightarrow j}$  is defined as:

$$e_{i \rightarrow j} = \begin{cases} l & \text{if } \text{pos}_j - \text{pos}_i > l, \\ -l & \text{elif } \text{pos}_j - \text{pos}_i < -l, \\ \text{pos}_j - \text{pos}_i & \text{otherwise;} \end{cases} \quad (1)$$



In addition to pharma responsibilities, the pharma offered general medical advice and a range of services that are now performed solely by other specialist practitioners, such as surgery and midwifery.

Figure 2: A partial view of the SRL graph corresponding to the given sentence. **Blue**: sentence root; **red**: predicates/verbs; **green**: arguments; **yellow**: modifiers.

in which,  $\text{pos}_i$  and  $\text{pos}_j$  indicate the two tokens' position indices in  $s$ . Therefore, the total number of relations  $\mathcal{R} = 2l + 1$ , the set of relations consists of all integers from  $-l$  to  $l$ . We also connect tokens with themselves via self-connections ( $e_{i \rightarrow i} = 0$ ) to facilitate message passing in GNNs.

### 3.1.3 Semantic Role Labeling (Parser-based)

$G \in \{0, 1\}^{\mathcal{R} \times \mathcal{N} \times \mathcal{N}}$ ,  $\mathcal{V}$ : chunks returned by SRL,  
 $\mathcal{E}$ : semantic role labels of the chunks.

Similar to recent approaches designed for static MRC tasks (Zheng and Kordjamshidi, 2020; Zhang et al., 2020), we investigate building knowledge graphs via Semantic Role Labeling (SRL). An SRL system (Shi and Lin, 2019) can detect the arguments associated with each of the predicates (or verbs) within a sentence and how they are classified into specific roles. This property is essential in tackling MRC tasks, especially for extractive QA datasets, where answers are typically short chunks of text (e.g., entities), such chunks can often be identified as arguments by an SRL parser. Via SRL, we can easily know how arguments interact with each other, further connecting common chunks detected in multiple sentences can produce an argument flow that helps understanding the paragraph.

In our SRL graphs<sup>3</sup>, we use chunks that are identified as predicates, arguments, and modifiers as nodes. We use their semantic role labels w.r.t. their corresponding predicates as relations. For longer sentences, an SRL system often returns multiple predicate-centric graphs. We define a special **ROOT** node for each sentence, and connect it with all the predicates within the sentence, with a special relation **ROOT-VERB**. We connect sentences by con-

<sup>3</sup>SRL is done with AllenNLP Toolkit (<https://demo.allennlp.org/semantic-role-labeling>).

necting their root nodes using a relation **ROOT-ROOT**. An example of the SRL graph is shown in Figure 2.

In order to facilitate models such as GNNs (e.g., easier message passing, denser signals), we define a set of reversed relations in SRL graphs. For instance, in Figure 2,  $e_{\text{performed} \rightarrow \text{services}}$  is **ARG1**, we use an **ARG1-rev** relation as  $e_{\text{services} \rightarrow \text{performed}}$  in our experiments. However, we ignore the reversed relations in Figure 2 for simplicity.

### 3.1.4 Continuous Belief (Trainable model)

$G \in [-1, 1]^{\mathcal{R} \times \mathcal{N} \times \mathcal{N}}$ ,  $\mathcal{V}$ : concepts determined by agent,  
 $\mathcal{E}$ : relations determined by agent.

In addition to rule- and parser-based graph updaters, we also investigate a data-driven approach. Inspired by Adhikari et al. (2020), we use self-supervised learning technique to pre-train a graph updater that can maintain a *continuous* belief graph for the iMRC task. Specifically, given the new observation  $o_t$ , we use a neural network to modify graph  $G_{t-1}$  from previous step, to produce a new graph  $G_t$ . Without the need of manually defining and hand-crafting specific graph structures — which may inject unnecessary prior into the agent — we assume that as long as a learned graph  $G_t$  can be used to reconstruct  $o_t$ , the graph should have contained useful information of the text.

However, learning to reconstruct  $o_t$  in a word-by-word fashion requires the model to learn features that are less useful in the iMRC task. Therefore, we adopt a contrastive representation learning strategy to approximate the reconstruction. Specifically, we train a discriminator  $D$  that differentiates between true  $o_t$  (positive samples) and a “corrupted” version of them  $\tilde{o}_t$  (negative samples), conditioned on  $G_t$ . This relieves the model of the burden to learn syntactical features (NLG), so that it can focus more on the semantic side instead. We use a standard noise-contrastive objective to minimize the binary cross-entropy (BCE) loss (Veličković et al., 2019):

$$\mathcal{L} = \frac{1}{K} \sum_{t=1}^K (\mathbb{E}_o[\log D(h_{o_t}, h_{G_t})] + \mathbb{E}_{\tilde{o}}[\log(1 - D(h_{\tilde{o}_t}, h_{G_t}))]), \quad (2)$$

where  $K$  is the number of sentences in a SQuAD paragraph. To facilitate this pre-training, we utilize an online Wikipedia dump (Wilson, 2013). We remove all articles that appear in the SQuAD dataset, and use the rest as our negative sample collection<sup>4</sup>.

<sup>4</sup>We apply the filtering to prevent the pre-trained graph updater from “memorizing” text that may appear in SQuAD validation set



The graph updater is then trained to generate graphs  $G_t$  that can be used to differentiate between 1) sentences within current document and 2) sentences sampled from another Wikipedia article.

Note that  $G_t$  is not explicitly grounded to any ground-truth graphs. Instead, they are essentially latent recurrent state representations, encoding information the agent has seen so far. Therefore, the nodes and relations in these graphs are determined by the agent itself in a data-driven manner. As a result, the adjacency tensors in these graphs are real-valued. We provide more details of this graph updater in Appendix A.4.

### 3.2 Encoding Graph Representations

We adopt a multi-layer relational graph convolutional network (R-GCN) (Schlichtkrull et al., 2018; Adhikari et al., 2020) as our graph encoder. Specifically, at the  $l$ -th layer of the R-GCN, for each node  $i \in \mathcal{V}$ , given the set of its neighbor nodes  $\mathcal{V}_i^e \in \mathcal{V}$  under relation  $e \in \mathcal{E}$ , the R-GCN computes:

$$\tilde{h}_i = \sigma \left( \sum_{e \in \mathcal{E}} \sum_{j \in \mathcal{V}_i^e} W_e^l [h_j^l; \text{Emb}_e] + W_0^l [h_i^l; \text{Emb}_e] \right), \quad (3)$$

where  $W_e^l$  and  $W_0^l$  are trainable parameters. When the graph is discrete (i.e., word co-occurrence, relative position, SRL), we use ReLU as the activation function  $\sigma$ ; when the graph is continuous (i.e., continuous belief), we use Tanh function as  $\sigma$  to stabilize the model.

When the labels of graph nodes consist of tokens, we integrate their word representations into graph computation. Specifically, for a node  $i$ , we use the concatenation of a randomly initialized node embedding vector and the averaged word embeddings of node label as the initial input  $h_i^0$ . Similarly, for each relation  $e$ ,  $\text{Emb}_e$  is the concatenation of a randomly initialized relation embedding vector and the averaged word embeddings of  $e$ 's label.

We utilize highway connections (Srivastava et al., 2015) between R-GCN layers:

$$\begin{aligned} g &= \text{Sigmoid}(W_{\text{hw}}(\tilde{h}_i)), \\ h_i^{l+1} &= g \odot \tilde{h}_i + (1 - g) \odot h_i^l, \end{aligned} \quad (4)$$

where  $\odot$  indicates element-wise multiplication,  $W_{\text{hw}}$  is a linear layer. We denote the final output of the R-GCN as  $h_{G_t} \in \mathbb{R}^{\mathcal{N} \times H}$ , where  $\mathcal{N}$  is the number of nodes in the graph,  $H$  is hyperparameter.

### 3.3 Aggregating Multiple Modalities

Following Yuan et al. (2020), we utilize the context-query attention mechanism (Yu et al., 2018) to aggregate multiple representations. The inputs to a context-query attention layer are typically two sequences of representation vectors (e.g., sequence of tokens for text, sequence of nodes for graphs). The attention computes element-wise similarity scores between the two inputs, then each element in one input can be represented by the weighted sum of the other input, and vice versa.

As shown in Figure 1 (right), we stack another context-query attention layer on top of the encoder used in iMRC, to aggregate the text representation (which encodes information in  $o_t$  and  $q$ ) with graph representation. We denote the output from the second attention layer as  $h_{og} \in \mathbb{R}^{L_{ot} \times H}$ , where  $L_{ot}$  is the length of  $o_t$ ,  $H$  is hyperparameter.

Although all the four graph types we investigate are updated dynamically, they can only represent the agent's belief of the *current* state  $s_t$ . There are clearly some information hard to be represented in the graphs, such as how did an agent navigate to the current sentence (i.e., the trajectories). We thus leverage a recurrent neural network to incorporate history information into encoder's output representations. Specifically, we use a GRU (Cho et al., 2014) as the recurrent component:

$$\begin{aligned} h_{\text{inp}} &= \text{MaskedMean}(h_{og}), \\ M_t &= \text{GRU}(h_{\text{inp}}, M_{t-1}), \end{aligned} \quad (5)$$

in which,  $h_{\text{inp}} \in \mathbb{R}^H$ .  $M_{t-1}$  is the output of the GRU cell at game step  $t - 1$ . As shown in Figure 1 (left), the output of encoder,  $M_t$ , is then used to both generating actions during information gathering phase, as well as extracting answers during question answering phase, this procedure exactly follows the iMRC pipeline.

## 4 Experiments and Results

We conduct experiments on the iSQuAD dataset (Yuan et al., 2020) to answer three key questions:

- Q1: Do graph representations help agents achieving better performance? In particular, among the four graph types, which of them provides the most performance boost?
- Q2: Do graph representations remain helpful in settings where multiple memory slots (observation queues) are available?

#Mem Slot	Agent	Easy Mode			Hard Mode			%RI	Reference
		$q$	$q + o_t$	$vocab$	$q$	$q + o_t$	$vocab$		
1	iMRC (Yuan et al., 2020)	0.575	0.579	0.583	0.524	0.357	<b>0.264</b>	–	§ 4.A1
	Ours (co-occur)	0.632	0.624	0.635	0.582	0.426	0.258	9.16	
	Ours (rel. pos.)	<b>0.634</b>	0.634	<b>0.642</b>	0.562	<b>0.440</b>	0.250	9.18	
	Ours (SRL)	0.616	<b>0.641</b>	0.638	<b>0.603</b>	0.434	0.253	<b>9.98</b>	
	Ours (cont.)	0.617	0.628	0.616	0.597	0.436	0.257	9.14	
	Ours (ensemble)	0.677	0.691	0.686	0.627	0.472	0.276	18.53	
3	iMRC (Yuan et al., 2020)	0.637	0.651	0.624	0.524	0.362	0.261	–	§ 4.A2
	Ours (co-occur)	0.674	0.665	<b>0.675</b>	0.605	<b>0.446</b>	0.260	<b>9.06</b>	
	Ours (rel. pos.)	0.677	0.665	0.664	<b>0.615</b>	0.438	0.257	8.66	
	Ours (SRL)	<b>0.681</b>	<b>0.678</b>	0.654	0.600	0.440	0.258	8.45	
	Ours (cont.)	0.676	0.642	0.662	0.592	0.426	<b>0.282</b>	8.26	
	Ours (ensemble)	0.714	0.713	0.701	0.650	0.471	0.278	15.80	
5	iMRC (Yuan et al., 2020)	0.666	0.656	0.661	0.551	0.364	0.218	–	§ 4.A2
	Ours (co-occur)	0.680	0.670	0.665	0.628	0.444	<b>0.258</b>	9.84	
	Ours (rel. pos.)	0.686	0.677	0.665	0.622	0.446	0.253	9.76	
	Ours (SRL)	0.675	0.680	0.680	0.609	0.441	0.257	9.56	
	Ours (cont.)	<b>0.699</b>	<b>0.693</b>	<b>0.696</b>	<b>0.629</b>	<b>0.455</b>	0.257	<b>12.19</b>	
	Ours (ensemble)	0.736	0.733	0.725	0.665	0.484	0.277	18.79	

Table 1: Testing  $F_1$  scores and the relative improvement %RI (averaged over six settings in a row). Best **single agent** scores within each setting are highlighted with **boldface**, scores better than iMRC are shaded in yellow.

- Q3: If graph representations are great, can we get rid of the text modality?

**Experiment Setup:** The iSQuAD dataset (Yuan et al., 2020) is an interactive version of the SQuAD dataset (Rajpurkar et al., 2016), which consists of 82k/5k/10k environments for training, validation, and testing. As described in Section 2, iSQuAD contains two difficulty levels and three finer-grained **QUERY** type settings, all of which influence an RL agent’s action space. The environment provides an observation queue with  $k$  memory slots depending on different configurations, where  $k \in \{1, 3, 5\}$ . The observation queue stores the  $k$  most recent observation sentences to alleviate difficulties caused by partial observability. Note in configuration where  $k = 1$ , there is no history information stored.

Inherited from the original SQuAD dataset, an agent is evaluated by the  $F_1$  score between its predicted answer and the ground-truth answers. We compare our agents equipped with graph representations against iMRC scores reported in (Yuan et al., 2020), specially, we also report an agent  $m$ ’s relative improvement over iMRC:

$$\%RI = (F_1^m - F_1^{iMRC}) / F_1^{iMRC} \times 100.0. \quad (6)$$

For all experiment settings, we train the agent with three different random seeds. We compute an agent’s test score using the model checkpoint that achieves the best validation score.

### A1: Graph representations indeed help, and ensemble is an useful strategy.

Intuitively, a dynamically maintained graph can serve as an agent’s episodic memory. Therefore, the less information is provided by the environment, the more useful the graphs can be. We first investigate our graph aided agent’s performance on the game configuration where only single memory slot is available. This is arguably the most difficult configuration in iMRC, where any valid action can lead to a completely different observation (a new sentence). As a result, agents needs to rely on its own memory mechanism.

As shown in Table 1 (#Mem Slot = 1), our agent outperforms iMRC in most of the settings by a noticeable margin. We observe that the improvement brought by graph representations is consistent across the four graph types. All of the four graph types provide over 9% of average relative improvement over iMRC. Among the four graph types, relative position graph and SRL graph seem to show advantage over the other two types, but this trend is not as significant.

Following standard strategy of model ensemble in MRC works, we test the ensemble of the four graphs. Specifically, taking four individual agents, each trained with its corresponding graph types, we mix their decisions during test. During the information gathering phase, we sum up the four agents’ output probabilities, including the probabilities over action words (i.e., **previous**, **next**, **ctrl+f**, **stop**) and the probabilities over the **QUERY** tokens.

Text Only (Yuan et al., 2020)				
0.575				
Graph Only				
co-occur	rel. pos.	SRL	cont.	ensemble
0.543	0.528	0.398	0.308	0.534
Text + Graph				
co-occur	rel. pos.	SRL	cont.	ensemble
0.632	0.634	0.616	0.617	0.677

Table 2: Testing  $F_1$  with different input modalities. Reference: § 4.A3.

The four agents consequently take the action with the max summed probabilities to keep interacting with the environment. During the question answering phase, we also sum up the output probabilities (over tokens in the sentence where the agents stop), and generate answers accordingly. Surprisingly, we find that the ensemble greatly boosts agent’s performance. As shown in Table 1 (#Mem Slot = 1), the ensemble agent nearly doubles our agent’s relative improvement over iMRC. It is also worth noting that with ensembling, our agent achieved better score than iMRC in the Hard Mode + *vocab* setting, which all the individual agent fail to outperform the baseline. This observation aligns with our motivation that the four types of graphs capture different aspects of the information and thus may be complementary to each other.

## A2: Graph representations remain helpful even with explicit memories.

As mentioned above, in some configurations, the iSQuAD environment provides an observation queue that caches most recent few observations as an explicit memory mechanism. A natural thing to explore is that if the advantages of equipping graph representations tend to diminish when the partial observability of the environments decreases.

We train and test our agent using iSQuAD’s configurations where 3 or 5 memory slots are available (i.e., at game step  $t$ , the input  $o_t$  to the agent is the concatenation of the most recent 3 or 5 sentences it has seen). From Table 1, we observe that the previously observed trends are consistent across different memory slot number configurations. Particularly, in the settings with 3 or 5 memory slots, our single agents equipped with graph representations can outperform iMRC in most of the settings. All graph types provide a greater than %8 of averaged relative improvements. Again, the ensemble agent nearly doubles single agents’ relative improvement over iMRC. In the setting with 5 memory slots, we

observe that the continuous belief graph is consistently outperforming its counterparts, which provides a %12.19 of relative improvement.

Given the observation that graph representations seem still helpful even with explicit memories, we further compare graph as memory mechanism against the explicit memory slots provided by iSQuAD environments. Comparing our best graph aided agent (receiving single sentence as input) against iMRC (receiving 3 and 5 sentences as input), we find our agent achieves a %13.03 and %13.47 of relative improvements over iMRC. This suggests that the design of the memory mechanism plays a big role in the interactive reading comprehension tasks. Although the concatenation of memory slots may provide as much amount of information, the inductive bias of graph representations are stronger.

## A3: Text modality is necessary.

Based on our findings in previous subsections, we further investigate whether the dynamic graphs can replace the text modality. We conduct a set of ablation experiments on the Easy Mode +  $q$  games, with single memory slot. Specifically, at every game step  $t$ , given the new observation  $o_t$ , we use the graph updater (described in Section 3.1) to generate graph representations  $G_t$ . Encoded by the graph encoder (described in Section 3.2), we directly aggregate the graph encoding with the question representations for further computations. In this way, the observation sentence  $o_t$  is only used to build the graph, without serving as a direct input modality to the agent.<sup>5</sup>

After training and testing such graph-only variants of our agent, we compare them against the text-only version (Yuan et al., 2020) and our full agent with both input modalities in Table 2. We observe that the graph-only agent fails to outperform the text-only baseline with any of the graph types, even with the ensemble of them. This suggests that even though the text and graph modalities may contain redundant information (because the graphs are generated from the text), they represent the information complementarily in some sense. We suspect the attention mechanism integrating text representations and graph representations (described in Section 3.3) may have contributed to the improvement of the full agent. For instance, the

<sup>5</sup>Note  $o_t$  is absent only during the information gathering phase, due to the extractive design of the question answerer.

#Mem	Agent	co-occur	rel. pos.	SRL	cont.	ensemble
1	Ours	9.16	9.18	9.98	9.14	18.53
	w/o RNN	4.73	4.27	5.50	5.23	12.29
3	Ours	9.06	8.66	8.45	8.26	15.80
	w/o RNN	5.89	7.48	2.91	5.43	14.15
5	Ours	9.84	9.76	9.56	12.19	18.79
	w/o RNN	6.51	5.97	6.29	6.02	14.99

Table 3: Averaged %RI over iMRC, comparing full agent with variants without RNN in encoder. Reference: § 4. Additional Results.

agent may have learned to focus on certain sub-graph conditioned on tokens in  $o_t$ , and vice versa.

### Additional Results and Discussion

As described in Section 3.3, our agent utilizes a recurrent component to be aware of history information, this component is absent in the original iMRC architecture. Therefore, it is important to make sure the performance improvement shown in previous subsections are not solely caused by the RNN. We conduct a set of experiments, with the RNN layer disabled (i.e., the output of attention layer becomes  $M_t$  in Figure 1 right).

We show results of these experiments in Table 3, due to space limitation, we only show the averaged relative improvement over iMRC, readers can find full results in Appendix C. Overall, graph representations contribute more to the improvement (for single agents, more than %5 on average), this is especially clear for the ensemble agents, where even without RNN, agents can sometimes achieve very close performance with the full agent. However, the effect of the RNN is non-negligible. This again emphasizes the importance of memory mechanism in interactive reading comprehension tasks. From our finding, multiple distinct memory mechanisms (i.e., memory slots, graphs, RNN cells) do not seem redundant, rather, they work cooperatively to produce a better score than solely using any of them.

It is noticeable in Table 1 that all agents performs poorly on Hard Mode + *vocab* games. This reveals limitations of RL-based algorithm (such as deep Q-learning we use in this work) — when the action space is extremely large, the agent has near-zero probability to experience a trajectory that leads to any positive reward, and thus struggles to learn useful strategies. This can potentially be mitigated by pre-training the agent with an easier setting then fine-tune in the difficult setting so that the agent has higher probability to experience good trajectories

to start with.

A recent work (Guo et al., 2021) propose to facilitate RL learning in tasks with huge action spaces (e.g., natural language generation) using Path Consistency Learning (PCL). Their PCL-based training method can update Q-values of all actions (tokens in vocabulary) at once, as opposed to only update the selected action (one token) in vanilla Q-Learning. This can potentially enable iMRC agents to perform in a more natural and generic manner, for instance, to **Ctrl+F** multi-word expressions as **QUERY**.

Due to space limitation, we report detailed agent structure, more results, and implementation details in Appendices.

## 5 Related Work

MRC has become an ever-growing area in the past decade, especially since the success of deep neural models. Like an adversarial game, researchers release new datasets (Hill et al., 2015; Chen et al., 2016; Rajpurkar et al., 2016; Trischler et al., 2016a; Nguyen et al., 2016; Reddy et al., 2018; Yang et al., 2018a; Choi et al., 2018; Clark et al., 2020) and novel models (Trischler et al., 2016b; Wang and Jiang, 2016; Seo et al., 2016; Wang et al., 2017; Huang et al., 2018) one after another. Since the flourishing of large scale pre-trained language models such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) and XLNet (Yang et al., 2019), performance of neural models on MRC datasets have improved greatly.

While some researchers believe models have achieved human-level performance, others argue that there have been biases or trivial cues injected into MRC datasets unconsciously (Agrawal et al., 2016; Weissenborn et al., 2017; Mudrakarta et al., 2018; Sugawara et al., 2018; Niven and Kao, 2019; Sen and Saffari, 2020). These biases may cause models to learn shallow pattern matching, rather than deep understanding and reasoning skills.

iMRC (Yuan et al., 2020) is a line of research that assumes partial observability and insufficient information. To answer a question, models have to actively collect necessary information by interacting with the environment. The iMRC paradigm can be described naturally within the RL framework, and thus it shares interests with video games (Badia et al., 2020), text-based games (Ammanabrolu and Riedl, 2019; Adhikari et al., 2020) and navigation (Anderson et al., 2018; Shridhar et al., 2020).



Graph construction is also a thriving direction lies at the intersection of multiple areas such as information extraction (Angeli et al., 2015), knowledge base (Shin et al., 2015), logical reasoning (Sinha et al., 2019) and representation learning (Kipf et al., 2020). Leveraging automatically constructed graph representations in static MRC has been shown effective, researchers use a wide range of linguistic features to help constructing graphs. Dhingra et al. (2018); Song et al. (2018) build graphs use coreference relations, De Cao et al. (2019); Qiu et al. (2019); Tu et al. (2020) leverage mentions of entities and Zheng and Kordjamshidi (2020) build SRL graphs using parsers. In the context of RL, prior work have also shown that constructing graph representations from other modalities can be helpful to solve tasks in interactive environments (Johnson, 2017; Yang et al., 2018b; Ammanabrolu and Riedl, 2019; Adhikari et al., 2020).

## 6 Broader Impact

Our work is a proof-of-concept study, we use a relatively simple and restricted (in terms of both observations and actions) QA dataset, iSQuAD, for both training and evaluation. Although the current version of our work might have limited consequences for society, we believe that taking a broader view of our work can be beneficial by preventing our future research from causing potential social and ethical concerns.

Similar to many RL-based systems, the information gathering module of our agent is optimized solely on its performance w.r.t. the final metric, without much constraints on its behavior at each game step. This can potentially make the system vulnerable since the RL agent may develop undesirable strategies that optimize the final metric.

In our current setting, the action space of the information gathering module is restricted (see Section 2). However, if we consider a more general setting, e.g., to equip the agent with a larger action space by allowing it to generate a sequence of tokens as the **QUERY** to the **Ctrl+F** action, we have to be extra careful about the aforementioned side effects caused by RL training. For instance, the agent may develop unfavorable behaviors such as forgetting proper syntax, abusing certain pronouns, to optimize its final rewards.

## 7 Conclusion

We explore to leverage graph representations in the challenging iMRC tasks. We investigate different categories of graph structures that can capture text information at various levels. We describe methods that dynamically generate the graphs during information gathering. Experiment results show that graph representations provide consistent improvement across settings. This evinces our hypothesis that graph representations are proper inductive biases in iMRC.

## Acknowledgments

We thank Marc-Alexandre Côté, Jie Fu and Tong Wang for the helpful discussions about this work. We also thank the anonymous EMNLP reviewers and area chairs for their helpful feedback and suggestions.

## References

- Ashutosh Adhikari, Xingdi Yuan, Marc-Alexandre Côté, Mikuláš Zelinka, Marc-Antoine Rondeau, Romain Laroche, Pascal Poupart, Jian Tang, Adam Trischler, and William L. Hamilton. 2020. Learning dynamic belief graphs to generalize on text-based games. In *NeurIPS*.
- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. 2016. *Analyzing the behavior of visual question answering models*. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1955–1960, Austin, Texas. Association for Computational Linguistics.
- Prithviraj Ammanabrolu and Mark Riedl. 2019. *Playing text-adventure games with graph-based deep reinforcement learning*. 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)*, pages 3557–3565, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. *Leveraging linguistic structure for open domain information extraction*. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages

- 344–354, Beijing, China. Association for Computational Linguistics.
- Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.
- Philip Bachman, Alessandro Sordoni, and Adam Trischler. 2016. Towards information-seeking agents. *CoRR*, abs/1612.02605.
- Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturovski, Pablo Sprechmann, Alex Vitvitskiy, Zhao-han Daniel Guo, and Charles Blundell. 2020. Agent57: Outperforming the Atari human benchmark. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 507–517. PMLR.
- Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. A thorough examination of the CNN/Daily Mail reading comprehension task. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2358–2367, Berlin, Germany. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium. Association for Computational Linguistics.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2019. Question answering by reasoning across documents with graph convolutional networks. 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)*, pages 2306–2317, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2018. 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)*, pages 42–48, New Orleans, Louisiana. Association for Computational Linguistics.
- James Ferguson, Matt Gardner, Hannaneh Hajishirzi, Tushar Khot, and Pradeep Dasigi. 2020. IIRC: A dataset of incomplete information reading comprehension questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1137–1147, Online. Association for Computational Linguistics.
- Han Guo, Bowen Tan, Zhengzhong Liu, Eric P Xing, and Zhiting Hu. 2021. Text generation with efficient (soft) q-learning. *arXiv preprint arXiv:2106.07704*.
- Matthew Hausknecht and Peter Stone. 2015. Deep recurrent q-learning for partially observable mdp. *arXiv preprint arXiv:1507.06527*.
- Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2015. The goldilocks principle: Reading children’s books with explicit memory representations. *arXiv preprint arXiv:1511.02301*.
- Hsin-Yuan Huang, Chenguang Zhu, Yelong Shen, and Weizhu Chen. 2018. Fusionnet: Fusing via fully-aware attention with application to machine comprehension. In *International Conference on Learning Representations*.
- Daniel D Johnson. 2017. Learning graphical state transitions. In *International Conference on Learning Representations (ICLR)*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Thomas Kipf, Elise van der Pol, and Max Welling. 2020. Contrastive learning of structured world models. In *International Conference on Learning Representations*.

- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2020. Question and answer test-train overlap in open-domain question answering datasets. *arXiv preprint arXiv:2008.02637*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. [Did the model understand the question?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1896–1906, Melbourne, Australia. Association for Computational Linguistics.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. [MS MARCO: A human generated machine reading comprehension dataset](#). *CoRR*, abs/1611.09268.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Lin Qiu, Yunxuan Xiao, Yanru Qu, Hao Zhou, Lei Li, Weinan Zhang, and Yong Yu. 2019. Dynamically fused graph network for multi-hop reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6140–6150.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2018. [Coqa: A conversational question answering challenge](#). *CoRR*, abs/1808.07042.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.
- Priyanka Sen and Amir Saffari. 2020. [What do models learn from question answering datasets?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2429–2438, Online. Association for Computational Linguistics.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Peng Shi and Jimmy Lin. 2019. Simple bert models for relation extraction and semantic role labeling. *arXiv preprint arXiv:1904.05255*.
- Jaeho Shin, Sen Wu, Feiran Wang, Christopher De Sa, Ce Zhang, and Christopher Ré. 2015. Incremental knowledge base construction using deepdive. In *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases*, volume 8, page 1310. NIH Public Access.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. [CLUTRR: A diagnostic benchmark for inductive reasoning from text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4506–4515, Hong Kong, China. Association for Computational Linguistics.
- Linfeng Song, Zhiguo Wang, Mo Yu, Yue Zhang, Radu Florian, and Daniel Gildea. 2018. Exploring graph-structured passage representation for multi-hop reading comprehension with graph neural networks. *arXiv preprint arXiv:1809.02040*.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. [Highway networks](#). *CoRR*, abs/1505.00387.
- Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. [What makes reading comprehension questions easier?](#) *CoRR*, abs/1808.09384.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2016a. [Newsqa: A machine comprehension dataset](#). *CoRR*, abs/1611.09830.



- Adam Trischler, Zheng Ye, Xingdi Yuan, Philip Bachman, Alessandro Sordoni, and Kaheer Suleman. 2016b. [Natural language comprehension with the EpiReader](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 128–137, Austin, Texas. Association for Computational Linguistics.
- Ming Tu, Kevin Huang, Guangtao Wang, Jing Huang, Xiaodong He, and Bowen Zhou. 2020. Select, answer and explain: Interpretable multi-hop reading comprehension over multiple documents. In *AAAI*, pages 9073–9080.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30*.
- Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2019. [Deep graph infomax](#). In *International Conference on Learning Representations*.
- Benyou Wang, Lifeng Shang, Christina Lioma, Xin Jiang, Hao Yang, Qun Liu, and Jakob Grue Simonsen. 2021. [On position embeddings in {bert}](#). In *International Conference on Learning Representations*.
- Shuohang Wang and Jing Jiang. 2016. [Machine comprehension using match-lstm and answer pointer](#). *CoRR*, abs/1608.07905.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. [Gated self-matching networks for reading comprehension and question answering](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 189–198, Vancouver, Canada. Association for Computational Linguistics.
- Dirk Weissenborn, Georg Wiese, and Laura Seiffe. 2017. [Making neural QA as simple as possible but not simpler](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 271–280, Vancouver, Canada. Association for Computational Linguistics.
- Benjamin Wilson. 2013. [The unknown perils of mining wikipedia](#).
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018a. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Zhilin Yang, Jake Zhao, Bhuwan Dhingra, Kaiming He, William W Cohen, Russ R Salakhutdinov, and Yann LeCun. 2018b. [Glomo: Unsupervised learning of transferable relational graphs](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8950–8961. Curran Associates, Inc.
- Adams Wei Yu, David Dohan, Quoc Le, Thang Luong, Rui Zhao, and Kai Chen. 2018. [Fast and accurate reading comprehension by combining self-attention and convolution](#). In *International Conference on Learning Representations*.
- Xingdi Yuan, Marc-Alexandre Côté, Alessandro Sordoni, Romain Laroché, Remi Tachet des Combes, Matthew Hausknecht, and Adam Trischler. 2018. Counting to explore and generalize in text-based games. *arXiv preprint arXiv:1806.11525*.
- Xingdi Yuan, Jie Fu, Marc-Alexandre Côté, Yi Tay, Chris Pal, and Adam Trischler. 2020. [Interactive machine comprehension with information seeking agents](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2325–2338, Online. Association for Computational Linguistics.
- Shuailiang Zhang, Hai Zhao, and Junru Zhou. 2020. Semantics-aware inferential network for natural language understanding. *arXiv preprint arXiv:2004.13338*.
- Chen Zhao, Chenyan Xiong, Jordan Boyd-Graber, and Hal Daumé III. 2021. Multi-step reasoning over unstructured text with beam dense retrieval. *arXiv preprint arXiv:2104.05883*.
- Chen Zheng and Parisa Kordjamshidi. 2020. [SRLGRN: Semantic role labeling graph reasoning network](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8881–8891, Online. Association for Computational Linguistics.



## Contents in Appendices:

- In Appendix A, we provide detailed information of our agent architecture.
- In Appendix B, we provide implementation details.
- In Appendix C, we provide the full set of our experiment results.

## A Details on Agent Structure

In this section, we provide detailed information of our agent. We will describe each of the modules shown in Figure 1. Some information here may be redundant with what we describe in Section 3, we repeat them here for reader’s convenience.

### Notations

We use *game step*  $t$  to denote one round of interaction between an agent with the iSQuAD environment. We use  $o_t$  to denote text observation at game step  $t$ , and  $q$  to denote question text. We use  $L$  to refer to a linear transformation, superscript of  $L$  denotes the activation function applied to the linear layer. Brackets  $[\cdot; \cdot]$  denote vector concatenation.

### A.1 Encoder

At a game step  $t$ , the encoder takes text observation (a sentence)  $o_t$ , the question  $q$ , and the graph  $G_t$  generated by the graph updater (if applicable) as input. It first converts each input into vector representations, then aggregates them using attention mechanism.

#### A.1.1 Text Encoder

We use a transformer-based encoder, which consists of an embedding layer and a transformer block (Vaswani et al., 2017). Specifically, embeddings are initialized by vectors extracted from a BERT model (Devlin et al., 2018) that is pre-trained on large corpus and fine-tuned on SQuAD<sup>6</sup>. The embedding size is 1024, they are fixed during training in all settings.

The transformer block consists of a stack of 4 convolutional layers, a self-attention layer, and a 2-layer MLP with a ReLU non-linear activation function in between. Within the block, each convolutional layer has 96 filters, with the kernel size of 7. In the self-attention layer, we use a block hidden

<sup>6</sup>We obtain the embeddings from HuggingFace (<https://huggingface.co/>), the BERT large model (uncased), whole word masking, fine-tuned on SQuAD.

size  $H$  of 96, as well as a single head attention mechanism. Layer normalization (Ba et al., 2016) is applied after each component inside the block. Following standard transformer training, we add positional embeddings into each block’s input.

At every game step  $t$ , we use the same text encoder to process  $o_t$  and  $q$ . The resulting representations are  $h_{o_t} \in \mathbb{R}^{L_{o_t} \times H}$  and  $h_q \in \mathbb{R}^{L_q \times H}$ , where  $L_{o_t}$  is the number of tokens in  $o_t$ ,  $L_q$  denotes the number of tokens in  $q$ ,  $H = 96$  is the hidden size.

#### A.1.2 Graph Encoder

We adopt the graph encoder from (Adhikari et al., 2020), which is a model based on R-GCN (Schlichtkrull et al., 2018). Specifically, at the  $l$ -th layer of the R-GCN, for each node  $i \in \mathcal{V}$ , given the set of its neighbor nodes  $\mathcal{V}_i^e \in \mathcal{V}$  under relation  $e \in \mathcal{E}$ , the R-GCN computes:

$$\tilde{h}_i = \sigma \left( \sum_{e \in \mathcal{E}} \sum_{j \in \mathcal{V}_i^e} W_e^l [h_j^l; \text{Emb}_e] + W_0^l [h_i^l; \text{Emb}_e] \right), \quad (7)$$

where  $W_e^l$  and  $W_0^l$  are trainable parameters. When the graph is discrete (i.e., word co-occurrence, relative position, SRL), we use ReLU as the activation function  $\sigma$ ; when the graph is continuous (i.e., continuous belief), we use Tanh function as  $\sigma$  to stabilize the model.

As the initial input  $h^0$  to the graph encoder, we concatenate a node embedding vector and the averaged word embeddings of node text (e.g., a word in word co-occurrence graph, a chunk of a sentence in SRL graph). Similarly, for each relation  $e$ ,  $\text{Emb}_e$  is the concatenation of a relation embedding vector and the averaged word embeddings of  $e$ ’s label. Both node embedding and relation embedding vectors are randomly initialized and trainable.

We utilize highway connections (Srivastava et al., 2015) between layers:

$$g = L^{\text{Sigmoid}}(\tilde{h}_i), \quad (8)$$
$$h_i^{l+1} = g \odot \tilde{h}_i + (1 - g) \odot h_i^l,$$

where  $\odot$  indicates element-wise multiplication.

We use a 3-layer graph encoder, with a hidden size  $H = 96$  in each layer. The node embedding size and relation embedding size are 100 and 32, respectively. The number of bases we use is 3. The final output of graph encoder is  $h_{G_t} \in \mathbb{R}^{\mathcal{N} \times H}$ , where  $\mathcal{N}$  is the number of nodes in the graph.

### A.1.3 Attention

**Context-Query Attention** To aggregate the question  $q$  with context comes from various modality (i.e., text and graph), we adopt the context-query attention layer from (Yu et al., 2018). We use a unified notion  $c$  to represent the context in the description of the context-query attention, the encoding of  $c$  is denoted as  $h_c \in \mathbb{R}^{L_c \times H}$ .

The attention layer first uses two MLPs to convert both  $h_c$  and  $h_q$  into the same space, the resulting tensors are denoted as  $h'_c \in \mathbb{R}^{L_c \times H}$  and  $h'_q \in \mathbb{R}^{L_q \times H}$ , in which  $H = 96$ . Then, a tri-linear similarity function is used to compute the similarities between each pair of  $h'_c$  and  $h'_q$  items:

$$S = W[h'_c; h'_q; h'_c \odot h'_q], \quad (9)$$

where  $W$  is trainable parameters with hidden size 96.

Softmax of the resulting similarity matrix  $S$  along both dimensions are computed, this produces  $S^A$  and  $S^B$ . Information in the two representations are then aggregated by:

$$\begin{aligned} h_{cq} &= [h'_c; P; h'_c \odot P; h'_c \odot Q], \\ P &= S_q h_q^{I\top}, \\ Q &= S_q S_c^\top h_c^{I\top}. \end{aligned} \quad (10)$$

Next, a linear transformation projects the aggregated representations to a space with size  $H = 96$ :

$$h_{cq} = L(h_{cq}). \quad (11)$$

Now,  $h_{cq} \in \mathbb{R}^{L_c \times H}$  is aggregated context-query representation.

**Context-Context Attention** Given the aggregated text-query representations  $h_{oq}$  and the aggregated graph-query representations  $h_{gq}$ , the context-context attention aims to merge them together and generate an overall representation that encodes all available information the agent has seen so far.

Specifically, the context-context attention is implemented as a stacked layers of transformer blocks. The structure of these transformer blocks is similar with the text encoder blocks, except we append an extra attention layer after its self-attention mechanism. This extra attention layer computes the attention between text-query representations with graph-query representations, followed with an extra layer normalization. In each block, we use a stack of 2 convolutional layers, each convolutional layer

has 94 filters with kernel size of 5. It is worth noting that this additional attention layer is performed only when graph representations are enabled, and is skipped otherwise. We stack 7 such transformer layers, they output  $h_{og} \in \mathbb{R}^{L_{o_t} \times H}$ , where  $L_{o_t}$  is the length of  $o_t$ ,  $H = 96$  is hidden size.

### A.1.4 Recurrent Component

As mentioned in Section 4, we have a setting where the encoder is recurrent, so that the agent can incorporate history information into its representations. In that specific setting, we use a GRU (Cho et al., 2014) as the recurrent component:

$$\begin{aligned} h_{\text{inp}} &= \text{MaskedMean}(h_{og}), \\ h_t &= \text{GRU}(h_{\text{inp}}, h_{t-1}), \end{aligned} \quad (12)$$

in which, the mean pooling is performed along the dimension of number of tokens, i.e.,  $h_{\text{inp}} \in \mathbb{R}^H$ .  $h_{t-1}$  is the output of the GRU cell at game step  $t - 1$ .

## A.2 Action Generator

Let  $M \in \mathbb{R}^H$  denote the output of the attention layers described above:

$$M = \begin{cases} h_t & \text{if recurrent,} \\ \text{MaskedMean}(h_{og}) & \text{otherwise.} \end{cases} \quad (13)$$

The action generator takes  $M$  as input and generates rankings for all possible actions. As defined in (Yuan et al., 2020), a **Ctrl+F** command is composed of two tokens (the token “**Ctrl+F**” and the **QUERY** token). Therefore, the action generator consists of three multi-layer perceptrons (MLPs):

$$\begin{aligned} h_{\text{shared}} &= L_{\text{shared}}^{\text{ReLU}}(M), \\ Q_{\text{action}} &= L_{\text{action}}(h_{\text{shared}}), \\ h_{\text{query}} &= L_{\text{query}}^{\text{Tanh}}(h_{\text{shared}}), \\ Q_{\text{query}} &= \text{Emb}(h_{\text{query}}). \end{aligned} \quad (14)$$

In which,  $Q_{\text{action}}$  and  $Q_{\text{query}}$  are Q-values of action token and **QUERY** token (when action token is “**Ctrl+F**”), respectively. The hidden size of  $L_{\text{shared}}$  is 150. The hidden size of  $L_{\text{action}}$  is either 2 (hard, only **Ctrl+F** and **stop** commands are allowed) or 4 (easy, **previous** and **next** are also allowed) depending on the game mode. We follow (Press and Wolf, 2017), tying the input embeddings and output embeddings. Specifically, a linear layer  $L_{\text{query}}$  followed by a Tanh activation projects  $h_{\text{shared}}$  into the same space as the embeddings (with dimensionality of 1024), then the pre-trained BERT embedding

matrix generates output logits  $Q_{\text{query}}$  (Q-values) where the output size is same as the vocabulary size.

Under different settings where the selection spaces of **QUERY** are specified, we apply different masks to  $Q_{\text{query}}$ . For instance, in the setting where the **QUERY** is a word selected from  $q + o_t$ , we use a mask which has same size as the vocabulary, where only tokens appear in either  $q$  and  $o_t$  are set to 1.

### A.3 Question Answerer

Whenever the action generator generates the command **stop**, or the agent has used up all its limit of moves, the information gathering phase terminates. At this point, the agent has to use its current internal representations to answer the question.

The question answer is a simple MLP-based layer. It takes  $h_{og} \in \mathbb{R}^{L_{o_t} \times H}$  as input, and generates a head distribution and a tail distribution over tokens in  $o_t$ :

$$\begin{aligned} h_{\text{head}} &= L_0^{\text{ReLU}}(h_{og}), \\ h_{\text{tail}} &= L_1^{\text{ReLU}}(h_{og}), \\ p_{\text{head}} &= L_2^{\text{Softmax}}(h_{\text{head}}), \\ p_{\text{tail}} &= L_3^{\text{Softmax}}(h_{\text{tail}}). \end{aligned} \quad (15)$$

### A.4 Graph Updater: Continuous Belief

Among all the four proposed graph categories, only the continuous belief graph is generated by the agent, and the graph updater is trained with a data-driven approach. Therefore, we describe the structure of this graph updater and the way we train it. Because a large portion of this module is adopted from (Adhikari et al., 2020), we provide a high level of the method and refer readers to (Adhikari et al., 2020) for detailed information.

We show the continuous belief graph updater training pipeline in Figure 3, it consists of two parts: the graph updater itself (red block on the left) and a decoder that helps to train the graph updater (green block on the right). As mentioned in Section 3.1.4, the idea is to train a graph updater that can modify and maintain a graph  $G_t$  using the text observation  $o_t$  and the graph from previous game step  $G_{t-1}$ . The graph  $G_t$  should contain sufficient information so that conditioned solely on  $G_t$ , a discriminator can differentiate true observation  $o_t$  from negative sample  $\tilde{o}_t$ .

In Figure 3, text encoder and graph encoder are similar modules as described in Appendix A.1. The  $f_{\Delta}$  function is a layer with attention mechanism

inside, it aggregates the text representations and graph representations. The  $f_{\Delta}$  function outputs a vector  $\Delta g_t$ , which represents the new information seen in the new observation  $o_t$ , compared to the graph at previous game step  $G_{t-1}$ .

The  $\oplus$  function is a graph operation function that produces the new belief representation  $h_t$  given  $h_{t-1}$  and  $\Delta g_t$ :

$$h_t = h_{t-1} \oplus \Delta g_t. \quad (16)$$

The graph operation function is implemented with a GRU (Cho et al., 2014). The function  $f_d$  is an MLP that decodes the recurrent state  $h_t$  into a real-valued adjacency tensor (i.e., the continuous belief graph  $G_t$ ).

At the decoder side (green block on the right), the graph representations of  $G_t$  are concatenated with both the text representations of  $o_t$  and  $\tilde{o}_t$ , the resulting vectors are fed into an MLP-based discriminator. The discriminator is trained with the standard binary cross-entropy (BCE) loss.

After pre-training, the graph updater (red block on the left) is fixed and plugged into our RL agent to produce continuous belief graphs.

## B Implementation Details

In this section, we provide hyperparameters and other implementation details.

For our full agent, we adopt the training procedure of DRQN (Hausknecht and Stone, 2015; Yuan et al., 2018) to train the agent. For the agent variants where recurrent component is absent, we use the training procedure of DQN (Mnih et al., 2013).

For all experiments, we use *Adam* (Kingma and Ba, 2015) as the optimizer. The learning rate is set to 0.00025 with a clip gradient norm of 5.

We use a prioritized replay buffer with memory size of 500,000, and a priority fraction of 0.5. During model update, we use a replay batch size of 64. We use a discount  $\gamma = 0.9$ . We use noisy nets, with a  $\sigma_0$  of 0.5. We update target network after every 1000 episodes. In DQN training, we sample the multi-step return  $n \sim \text{Uniform}[1, 3]$ . In DRQN training, every sampled data point in the replay batch is a sequence of 2 consecutive transitions. We use the first transition to estimate the recurrent states, and the second for updating the model parameters. We refer readers to Hessel et al. (2018) for more information about different components of DQN training.




Figure 3: Graph updater for continuous belief graphs.

We train all agents with 3 different random seeds. We choose the random seed which produces the best validation performance, and report its scores on the test set. To be comparable with iMRC, we also train our agents with 1 million episodes, each episode has a maximum number of steps 20.

We train all agents for 1 million episodes, this is the same number of episodes reported in (Yuan et al., 2020). Running speed of agents depend on the specific configuration, e.g., the type of graph equipped by an agent. On average, achieving best validation score takes an agent about 3 days on a single Nvidia P100 GPU.

## C Full Results

In Table 4,5,6, we provide full results on our experiments. Although the only metric to evaluate an agent’s performance on the iMRC task is the  $F_1$  score between the prediction with the ground-truth answers, in (Yuan et al., 2020), the authors also monitor agents’ sufficient information reward. Specifically, sufficient information rewards are binary rewards representing if the final observation (either the agent generates the `stop` action, or it has used up all its moves) contains the ground-truth answer as a sub-string. Intuitively, because of the extractive nature of the question answerer module, the agent can answer the question correctly if and only if it achieves a 1.0 sufficient information reward on a specific data point. We provide the sufficient information rewards of our agents in the full result tables, colored in blue.



Agent	Easy Mode			Hard Mode			%RI
	$q$	$q + o_t$	vocab	$q$	$q + o_t$	vocab	
iMRC (Yuan et al., 2020)	0.575 (0.747)	0.579 (0.739)	0.583 (0.753)	0.524 (0.684)	0.357 (0.477)	0.264 (0.363)	–
Ours (co-occur)	0.632 (0.779)	0.624 (0.763)	0.635 (0.770)	0.582 (0.724)	0.426 (0.533)	0.258 (0.338)	9.16 (3.40)
Ours (rel. pos.)	0.634 (0.776)	0.634 (0.774)	0.642 (0.779)	0.562 (0.696)	0.440 (0.553)	0.250 (0.338)	9.18 (3.81)
Ours (SRL)	0.616 (0.751)	0.641 (0.779)	0.638 (0.782)	0.603 (0.740)	0.434 (0.538)	0.253 (0.338)	9.98 (3.99)
Ours (cont.)	0.617 (0.757)	0.628 (0.763)	0.616 (0.758)	0.597 (0.744)	0.436 (0.542)	0.257 (0.338)	9.14 (3.46)
Ours (ensemble)	0.677 (0.789)	0.691 (0.799)	0.686 (0.795)	0.627 (0.735)	0.472 (0.555)	0.276 (0.338)	18.53 (6.04)
Ours (co-occur) w/o RNN	0.572 (0.708)	0.613 (0.753)	0.607 (0.744)	0.556 (0.689)	0.397 (0.495)	0.269 (0.359)	4.73 (-0.20)
Ours (rel. pos.) w/o RNN	0.574 (0.702)	0.621 (0.761)	0.608 (0.746)	0.540 (0.676)	0.409 (0.504)	0.255 (0.338)	4.27 (-1.05)
Ours (SRL) w/o RNN	0.559 (0.703)	0.628 (0.772)	0.631 (0.773)	0.540 (0.672)	0.411 (0.511)	0.266 (0.347)	5.50 (0.38)
Ours (cont.) w/o RNN	0.588 (0.724)	0.604 (0.739)	0.605 (0.747)	0.561 (0.693)	0.414 (0.513)	0.258 (0.338)	5.23 (-0.30)
Ours (ensemble) w/o RNN	0.598 (0.701)	0.659 (0.769)	0.664 (0.773)	0.593 (0.695)	0.441 (0.526)	0.278 (0.338)	12.29 (0.93)

Table 4: #Memory slot = 1. Testing  $F_1$  in **black** and sufficient information rewards in **blue**. %RI represents relative improvement over iMRC on corresponding metric, across settings.

Agent	Easy Mode			Hard Mode			%RI
	$q$	$q + o_t$	vocab	$q$	$q + o_t$	vocab	
iMRC (Yuan et al., 2020)	0.637 (0.738)	0.651 (0.734)	0.624 (0.738)	0.524 (0.740)	0.362 (0.729)	0.261 (0.719)	–
Ours (co-occur)	0.674 (0.863)	0.665 (0.859)	0.675 (0.872)	0.605 (0.780)	0.446 (0.585)	0.260 (0.338)	9.06 (3.52)
Ours (rel. pos.)	0.677 (0.874)	0.665 (0.850)	0.664 (0.877)	0.615 (0.789)	0.438 (0.575)	0.257 (0.338)	8.66 (3.54)
Ours (SRL)	0.681 (0.896)	0.678 (0.887)	0.654 (0.841)	0.600 (0.780)	0.440 (0.564)	0.258 (0.338)	8.45 (3.33)
Ours (cont.)	0.676 (0.870)	0.642 (0.850)	0.662 (0.864)	0.592 (0.746)	0.426 (0.544)	0.282 (0.389)	8.26 (3.49)
Ours (ensemble)	0.714 (0.864)	0.713 (0.867)	0.701 (0.851)	0.650 (0.783)	0.471 (0.567)	0.278 (0.338)	15.80 (2.75)
Ours (co-occur) w/o RNN	0.624 (0.810)	0.666 (0.870)	0.658 (0.860)	0.577 (0.732)	0.417 (0.523)	0.272 (0.362)	5.89 (0.34)
Ours (rel. pos.) w/o RNN	0.625 (0.840)	0.650 (0.838)	0.649 (0.835)	0.577 (0.741)	0.420 (0.524)	0.305 (0.386)	7.48 (1.19)
Ours (SRL) w/o RNN	0.593 (0.772)	0.639 (0.865)	0.636 (0.828)	0.568 (0.732)	0.400 (0.510)	0.275 (0.371)	2.91 (-1.18)
Ours (cont.) w/o RNN	0.641 (0.820)	0.645 (0.840)	0.667 (0.854)	0.576 (0.742)	0.420 (0.537)	0.261 (0.338)	5.43 (-0.57)
Ours (ensemble) w/o RNN	0.663 (0.803)	0.704 (0.859)	0.709 (0.858)	0.627 (0.748)	0.460 (0.538)	0.293 (0.350)	14.15 (0.31)

Table 5: #Memory slot = 3. Testing  $F_1$  in **black** and sufficient information rewards in **blue**. %RI represents relative improvement over iMRC on corresponding metric, across settings.

Agent	Easy Mode			Hard Mode			%RI
	$q$	$q + o_t$	vocab	$q$	$q + o_t$	vocab	
iMRC (Yuan et al., 2020)	0.666 (0.716)	0.656 (0.706)	0.661 (0.731)	0.551 (0.739)	0.364 (0.733)	0.218 (0.713)	–
Ours (co-occur)	0.680 (0.883)	0.670 (0.905)	0.665 (0.879)	0.628 (0.804)	0.444 (0.572)	0.258 (0.338)	9.84 (3.84)
Ours (rel. pos.)	0.686 (0.897)	0.677 (0.900)	0.665 (0.898)	0.622 (0.802)	0.446 (0.583)	0.253 (0.338)	9.76 (4.66)
Ours (SRL)	0.675 (0.884)	0.680 (0.902)	0.680 (0.920)	0.609 (0.767)	0.441 (0.563)	0.257 (0.338)	9.56 (3.43)
Ours (cont.)	0.699 (0.916)	0.693 (0.914)	0.696 (0.925)	0.629 (0.805)	0.455 (0.595)	0.257 (0.338)	12.19 (6.25)
Ours (ensemble)	0.736 (0.895)	0.733 (0.908)	0.725 (0.898)	0.665 (0.798)	0.484 (0.582)	0.277 (0.338)	18.79 (4.64)
Ours (co-occur) w/o RNN	0.631 (0.814)	0.674 (0.896)	0.682 (0.918)	0.586 (0.753)	0.412 (0.513)	0.259 (0.338)	6.51 (0.03)
Ours (rel. pos.) w/o RNN	0.636 (0.833)	0.682 (0.908)	0.669 (0.893)	0.572 (0.734)	0.412 (0.512)	0.258 (0.338)	5.97 (-0.29)
Ours (SRL) w/o RNN	0.619 (0.793)	0.683 (0.907)	0.635 (0.822)	0.571 (0.727)	0.434 (0.554)	0.265 (0.345)	6.29 (-0.72)
Ours (cont.) w/o RNN	0.638 (0.836)	0.664 (0.884)	0.653 (0.858)	0.575 (0.731)	0.427 (0.533)	0.259 (0.338)	6.02 (-0.72)
Ours (ensemble) w/o RNN	0.668 (0.803)	0.733 (0.905)	0.715 (0.879)	0.624 (0.742)	0.466 (0.546)	0.280 (0.338)	14.99 (0.16)

Table 6: #Memory slot = 5. Testing  $F_1$  in **black** and sufficient information rewards in **blue**. %RI represents relative improvement over iMRC on corresponding metric, across settings.