

Modeling Temporal-Modal Entity Graph for Procedural Multimodal Machine Comprehension

Huibin Zhang¹, Zhengkun Zhang¹, Yao Zhang¹, Jun Wang^{2*}, Yufan Li¹
Ning Jiang³, Xin Wei³, Zhenglu Yang^{1*}

¹TKLNDST, CS, Nankai University, China, ²Ludong University, China,

³Mashang Consumer Finance Co., Ltd., China

{zhanghuibin.cc, zhangzk2017, yaozhang, junwang, 1811486}
@mail.nankai.edu.cn, {ning.jiang02, xin.wei02}@msxf.com,
yangzl@nankai.edu.cn

Abstract

Procedural Multimodal Documents (PMDs) organize textual instructions and corresponding images step by step. Comprehending PMDs and inducing their representations for the downstream reasoning tasks is designated as Procedural MultiModal Machine Comprehension (M³C). In this study, we approach Procedural M³C at a fine-grained level (compared with existing explorations at a document or sentence level), that is, entity. With delicate consideration, we model entity both in its temporal and cross-modal relation and propose a novel Temporal-Modal Entity Graph (TMEG). Specifically, a heterogeneous graph structure is formulated to capture textual and visual entities and trace their temporal-modal evolution. In addition, a graph aggregation module is introduced to conduct graph encoding and reasoning. Comprehensive experiments across three Procedural M³C tasks are conducted on a traditional dataset RecipeQA and our new dataset CraftQA, which can better evaluate the generalization of TMEG.

1 Introduction

MultiModal Machine Comprehension (M³C) is a generalization of machine reading comprehension by introducing multimodality. Due to its differences from Visual Question Answering (VQA) (Antol et al., 2015) in the form of understanding multimodal contexts and conducting multimodal questions and answers, there has been a lot of attention in recent years devoted to this field. In this paper, we investigate a task that has been typical of M³C recently, named Procedural M³C, a task of reading comprehension of Procedural Multimodal Documents (PMDs).

As shown in Figure 1, a recipe that contains successive multimodal instructions is a typical PMD. Reading a recipe seems trivial for humans but is still complex for a machine reading comprehension

*Corresponding authors.

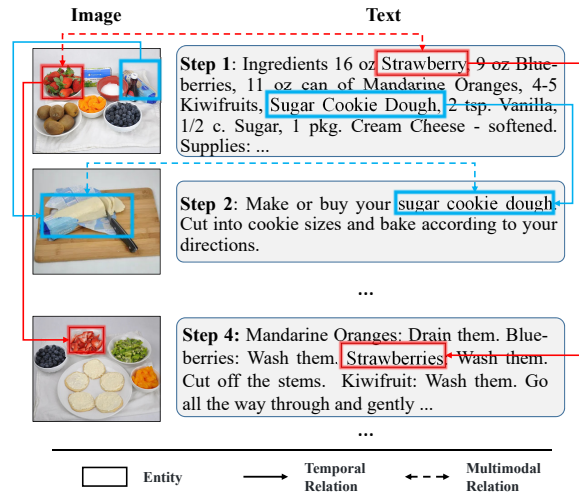


Figure 1: A PMD instance: a recipe about making mini fruit pizza. The entities and relations (i.e. temporal and multimodal) among them are highlighted.

system before it can comprehend both textual and visual contents and capture their relations.

Current Procedural M³C studies (Yagcioglu et al., 2018; Liu et al., 2020) comprehend PMDs by encoding text and images at each procedure step. These efforts, however, only scratch the surface and lack deep insight into the elementary unit of PMDs, that is, **entity**. From now on, we use entity to refer uniformly to entity in text and object in image. For instance, the recipe in Figure 1 involves multiple entities, i.e., Strawberry and Sugar Cookie Dough, etc. In this work, we target at approaching the Procedural M³C task at a fine-grained entity level.

We observe that a PMD essentially assembles an evolution process of entities and the relations between them. Specifically, the relation between entities can be summarized in the following two categories:

- **Temporal Relation.** The state of an entity may change as steps progress. Still looking at Figure 1, strawberries are complete at step 1 and

by step 4 they are washed and cut into pieces. We use temporal relation to depict the association between an entity’s changing visual signals in images or changing contexts in text.

- **Multimodal Relation.** An entity is naturally and powerfully associated with other entities within a single modality. Meanwhile, the cross-modal association of an entity is worth exploiting to contribute to distinct modality understanding. For example, the visual signal and context about **sugar cookie dough** can be interpreted by each other at step 2. We generalize the intra- and inter- modal associations of entities with the multimodal relation.

Based on the above observations, we believe that simultaneously modeling entities and the temporal and multimodal relations is a key challenge in understanding PMDs. Recent efforts (Amac et al., 2019; Huang et al., 2021) are devoted to encoding temporal relations of entities, while it neglects the multimodal relation. Heterogeneous graphs have become the preferred technology for representing, sharing, and fusing information to modern AI tasks, e.g., relation extraction (Christopoulou et al., 2019) and recommendation (Fan et al., 2019). Inspired by this, we construct a heterogeneous graph, with entities as nodes and relations as edges. Therefore, the research goal of the procedural M³C task will shift from understanding unstructured PMDs to learning structured graph representations.

In this work, we propose a novel **Temporal Modal Entity Graph** model, namely **TMEG**. Our model approaches Procedural M³C at a fine-grained entity level by constructing and learning a graph with entities and temporal and multimodal relations. Specifically, TMEG consists of the following components: 1) **Node Construction**, which extracts the token-level features in text and object-level features in images as the initial entity embeddings; 2) **Graph Construction**, which constructs the temporal, intra-modal, and cross-modal relations separately to form a unified graph; and 3) **Graph Aggregation**, which utilizes the graph-based multi-head attention mechanism to perform fusion operations on graphs to model the evolution of entities and relations. Finally, the graph representation is fed into a graph-based reasoning module to evaluate the model’s understanding ability.

In addition, in order to further advance the research of Procedural M³C, we release **CraftQA**, a multimodal semantically enriched dataset that con-

tains about 27k craft product-making tutorials and 46k question-answer pairs for evaluation. We evaluate three representative subtasks, i.e., visual cloze, visual coherence, and visual ordering, on **CraftQA** and a public dataset **RecipeQA** (Yagcioglu et al., 2018). The quantitative and qualitative results show the superiority of TMEG compared to the state-of-the-art methods on all three tasks. The main contributions of this paper can be summarized as follows:

- We innovatively study the Procedural M³C task at a fine-grained entity level. We comprehensively explore the relations between entities in both temporal and multimodal perspectives.
- We propose a Temporal-Modal Entity Graph model TMEG, which constructs a graph with entities as nodes and relations as edges, and then learns the graph representation to understand PMDs.
- We release a dataset **CraftQA**. The experimental results on **CraftQA** and **RecipeQA** show TMEG outperforms several state-of-the-art methods.

2 Related Work

Procedural Text Comprehension. Procedural text comprehension requires an accurate prediction for the state change and location information of each entity over successive steps. Several datasets have been proposed to evaluate procedural text comprehension, e.g., **Recipe** (Bosselut et al., 2018), **ProPara** (Mishra et al., 2019), and **OPENPI** (Tandon et al., 2020). To model entity evolution, **KG-MRC** (Das et al., 2019) constructs knowledge graphs via reading comprehension and uses them for entity location prediction. **DYNAPRO** (Amini et al., 2020) introduces a pre-trained language model to dynamically obtain the contextual embedding of procedural text and learn the attributes and transformations of entities. **ProGraph** (Zhong et al., 2020) enables state prediction from context by constructing a heterogeneous graph with various knowledge inputs. **Koala** (Zhang et al., 2021b) utilizes the external commonsense knowledge injection and data enhancement to reason the states and locations of entities. **TSLM** (Faghihi and Kordjamshidi, 2021) formulate comprehension task as a question answering problem and adapt pre-trained transformer-based language models on other QA benchmarks. **REAL** (Huang et al., 2021) builds

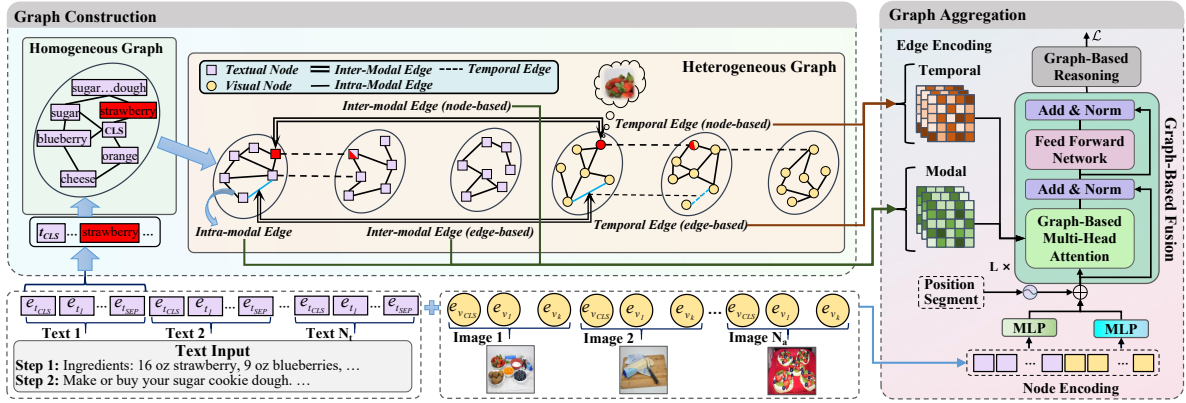


Figure 2: Overview of our proposed TMEG framework. We reuse the example in Figure 1. Initial nodes are generated from input text and images. Considering the temporal and cross-modal relation of entities, we apply various types of edges to construct a unified temporal-modal entity graph. By combining the encoding of nodes and edges, a graph aggregation module is designed to conduct graph encoding and reasoning.

a general framework to systematically model the entity, action, location by using a graph neural network.

Inspired by these previous works, we propose a temporal-modal entity graph model, which is designed with temporal encoding and modal encoding to model multiple types of entities.

Multimodal Graph. In recent multimodal research, graph structure has been utilized to model the semantic interaction between modalities. (Yin et al., 2020) propose a graph-based multimodal fusion encoder for neural machine translation, which converted sentence and image in a unified multimodal graph. (Khademi, 2020) convert image regions and the region grounded captions into graph structure and introduced graph memory networks for visual question answering. (Zhang et al., 2021a) propose a multimodal graph fusion approach for named entity recognition, which conducted graph encoding via multimodal semantic interaction. (Yang et al., 2021) focus on multimodal sentiment analysis and emotion recognition, which unified video, audio, and text modalities into an attention graph and learned the interaction through graph fusion, dynamic pruning, and the read-out technique.

In contrast to the above methods, we formulate our multimodal graph on temporal entities and successfully deploy it in Procedural M³C.

3 Proposed Method

In this section, we introduce: (1) problem definition of Procedural M³C in Section 3.1; (2) The homogeneous graph of each textual instruction (im-

age) and our TMEG in Section 3.2 and Section 3.3, respectively; and (3) graph aggregation module to conduct graph encoding and reasoning in Section 3.4. Figure 2 gives a high-level overview of TMEG.

3.1 Problem Definition

Here, we define the task of Procedural M³C, given:

- Context $\mathcal{S} = \{s_t\}_{t=1}^{N_t}$ in textual modality, which represents a series of coherent textual instructions to perform a specific skill or task (e.g., multiple steps to complete a recipe or a craft product).
- Question Q and Answer A , which is either a single image or a series of images in a reasoning task (e.g., visual cloze, visual coherence, or visual ordering).

Following (Liu et al., 2020), we combine the images contained in Q and A to form N_c candidate image sequences $\{a_1, \dots, a_j, \dots, a_{N_c}\}$. Let N_a be the length of the j -th candidate image sequence $a_j = \{I_{j,1}, \dots, I_{j,N_a}\}$. Take the visual cloze task as an example, we fill the placeholder of the question with candidate answers to form N_c image sequences with length $N_a = 4$. The model requires to select the most relevant candidate by calculating the similarity between text sequence $\mathcal{S} = \{s_t\}_{t=1}^{N_t}$ and each image sequence a_j .

3.2 Homogeneous Graph Construction

As shown in Figure 2, we first extract the tokens (objects) in text (image) as the initial nodes of homogeneous graph, respectively.

Textual Node. Let N_t be the number of textual instructions $\mathcal{S} = \{s_t\}_{t=1}^{N_t}$. First, each in-

struction s_t is tokenized into a token sequence $\{e_{t_{[\text{CLS}]}}\}, e_{t_1}, \dots, e_{t_{[\text{SEP}]}}\}$, where [CLS] and [SEP] are the special tokens introduced to mark the start and the end of each instruction. Then, we utilize an off-the-shelf POS tagger (Akbik et al., 2018) to identify all nouns and noun phrases in the token sequence. Finally, we concatenate all the token sequences of textual instructions \mathcal{S} and feed the token embedding into the graph aggregation module.

Visual Node. For each image sequence a_j , we employ a pre-trained Faster-RCNN to extract a set $\{e_{v_{[\text{CLS}]}}\}, e_{v_1}, \dots, e_{v_k}\}$ with k object features as visual tokens. Following (Messina et al., 2020; Dosovitskiy et al., 2021), we reserve [CLS] as the beginning token for each image whose final embedding is regarded as the representation of the whole image. The operation of visual node is in a similar manner as textual and any two nodes in the same instruction (image) are connected to construct a homogeneous graph.

3.3 Heterogeneous Graph Construction

Based on the homogeneous graph of each textual instruction (image), we introduce various types of edges to construct our heterogeneous graph TMEG.

3.3.1 Temporal Edge

It is essential to model the temporal evolution of entities for comprehending procedural content. Let us revisit the example in Figure 1. When a human reads step 4, the connection between entities (e.g., strawberries and oranges) and their descriptions in step 1 is naturally established.

We design the temporal edge to model the evolution of entities in text and image. It can be seen that the temporal edge describes the evolution at different steps. For the textual nodes, the same entity appearing in different steps are connected by a textual temporal edge (node-based). While for the visual nodes, we directly calculate the Euclidean Distance between object features due to the absence of accurate object detection. Following (Song et al., 2021), if the distance between node (object) i and node (object) j is less than a threshold λ_t , we treat them as the same object and connect node i to node j via a visual temporal edge (node-based). Meanwhile, we consider that there may also be temporal evolution for edges, such as changes in the relationship between entities. Therefore, we also introduce temporal edge (edge-based) to characterize the change of edges.

3.3.2 Modal Edge

As shown in Figure 1, the textual instruction of each image can be viewed as a noisy form of image annotation (Hessel et al., 2019). The association between image and sentence can be inferred through entity representations under different modalities. Correspondingly, we design the intra-modal edge and the inter-modal edge to represent the modal interactions. In Section 3.2, any two nodes in the same modality and the same instruction (image) are connected by an intra-modal edge. It is worth noting that for each instruction (image), the special [CLS] node is connected to all other nodes in order to aggregate graph-level features.

On the other hand, the textual node representing any entity and the corresponding visual node are connected by an inter-modal edge. We employ a visual grounding toolkit (Yang et al., 2019) to detect visual objects for each noun phrase. Specifically, we predict the bounding box corresponding to the text entity and compute the Intersection over Union (IoU) between all visual objects. If the IoU between the prediction box and the visual box exceeds a threshold λ_m , the textual node and the corresponding visual node are connected by an inter-modal edge (node-based). Similar to section 3.3.1, considering the influence of entity-relationship under different modalities, we also introduce inter-modal edge (edge-based) to characterize the interaction between edges.

3.4 Graph Aggregation

3.4.1 Node Encoding

As described in Section 3.2, we have obtained the embeddings of the textual tokens and visual objects. Similar to (Li et al., 2020), all embeddings are mapped to a set of initial node embeddings, and each node embedding is the sum of 1) a textual token embedding or visual object embedding; 2) a position embedding that identifies the position of the token or object in the image¹; and 3) a segment embedding generated from the step number in PMD which indicates different textual instructions or images.

3.4.2 Edge Encoding

To encode the structural information into TMEG, we consider the temporal encoding and the modal encoding separately. For any two nodes v_i and v_j in

¹We exploit the bounding box feature extracted by Faster-RCNN as the position embedding of the object.

TMEG, we construct two mappings: $\phi_t(v_i, v_j) \rightarrow \mathbb{R}$ and $\phi_m(v_i, v_j) \rightarrow \mathbb{R}$ which encode the temporal edge and the modal edge between them. The temporal encoding and the modal encoding of the total graph are fed into the graph-based aggregation module.

3.4.3 Graph-Based Fusion

As shown in the right part of Figure 2, we first introduce two multi-layer perceptrons (MLP) with Tanh activation function to project different node embeddings from two modalities into the same space. Then, we extend the VisualBERT (Li et al., 2020) to the graph-based fusion layer, which concatenates the node embeddings from MLPs as input and outputs their graph-based joint representations. Specifically, in each fusion layer, updating the hidden states of textual node and visual node mainly involve the following steps.

Firstly, we exploit a graph-based multi-head attention mechanism to generate contextual representations of nodes. Formally, the output of the h -th attention head in the $l - 1$ layer can be obtained as follows:

$$A(\mathbf{q}, \mathbf{k}, \mathbf{v})_j^{h,l-1} = \sum_{i=1}^N \mathbf{v}_i^h \left(\text{Softmax}(e_{i,j}^h) \right), \quad (1)$$

$$e_{i,j}^h = \frac{\mathbf{q}_j^h (\mathbf{k}_i^h)^T}{\sqrt{d}} + \mathbf{b}_{\phi_m(i,j)}^h + \mathbf{b}_{\phi_t(i,j)}^h, \quad (2)$$

where q, k, v are the query matrix, key matrix, and value matrix generated from the hidden state $H^{(l-1)}$ of nodes in the $l - 1$ layer. $\phi_t(i, j)$ and $\phi_m(i, j)$ denote the temporal encoding and the modal encoding of TMEG, which serve as bias terms in the attention module. It is worth noting that each head in the multi-head attention mechanism exhibits a broad range of behaviors (Vig, 2019); thus, we add different temporal encoding and modal encoding separately for each attention head. Meanwhile, in order to model the relationship of edges, the temporal encoding and the modal encoding are learned separately for each layer.

We concatenate the output of each head and pass them to a position-wise Feed Forward Networks (FFN) which is preceded and succeeded by residual connections and normalization layer (LN),

$$\begin{aligned} \hat{\mathbf{H}}^{(l)} &= \text{LN} \left(W[\mathbf{A}^1, \dots, \mathbf{A}^h] + \mathbf{H}^{(l-1)} \right), \\ \mathbf{H}^{(l)} &= \text{LN} \left(\text{FFN}(\hat{\mathbf{H}}^{(l)}) + \hat{\mathbf{H}}^{(l)} \right), \end{aligned} \quad (3)$$

where W is a learnable parameter and $[\]$ denotes

Algorithm 1: Graph Aggregation of TMEG

Input : The initial hidden states of TMEG \mathbf{H}^0 ; the graph-based fusion layer number N ; the temporal encoding Φ_t ; the modal encoding Φ_m ; the attention head \mathbf{A}

Output : The final hidden states \mathbf{H}^N

```

1 for  $l \leftarrow 1$  to  $N$  do
  // Graph-Based Fusion Layer
2   foreach attention head  $\mathbf{A}^{h,l}$  in layer  $l$  do
3     Generate  $\mathbf{q}, \mathbf{k}, \mathbf{v}$  from hidden states  $\mathbf{H}^{l-1}$ ;
4     Obtain edge encoding  $\mathbf{b}_{\phi_t(i,j)}^h$  and  $\mathbf{b}_{\phi_m(i,j)}^h$ ;
5     Calculate attention weight  $e_{i,j}^h$  (Eq. 2) with
        $\{\mathbf{q}, \mathbf{k}, \mathbf{b}_{\phi_t(i,j)}^h, \mathbf{b}_{\phi_m(i,j)}^h\}$ ;
6     Acquire output of head  $\mathbf{A}^{h,l}$  (Eq. 1) with
        $\{\mathbf{v}, e_{*,j}^h\}$ ;
7   end
8   Aggregate each attention head  $\mathbf{A}^h$ 
9   update
10   $\hat{\mathbf{H}}^{(l)} \leftarrow \text{LN}(W[\mathbf{A}^1, \dots, \mathbf{A}^h] + \mathbf{H}^{(l-1)})$ ;
11   $\mathbf{H}^{(l)} \leftarrow \text{LN}(\text{FFN}(\hat{\mathbf{H}}^{(l)}) + \hat{\mathbf{H}}^{(l)})$ ; (Eq. 3)
12 end
13 return  $\mathbf{H}^N$ 

```

the concatenation manipulation. Finally, based on TMEG, we stack multi-layer graph-based fusion layers to conduct graph encoding. Algorithm 1 shows the aggregation of TMEG in detail.

3.4.4 Graph-Based Reasoning

As mentioned in Section 3.2, we regard the hidden state of [CLS] as the representations of each instruction (image), where their final hidden states H^T and H^V are passed into the graph reasoning module for task completion.

Firstly, we leverage the one-to-one correspondence between instruction and image, e.g., each instruction has an image to visualize it (Alikhani et al., 2021). TMEG involves a Contrastive Coherence Loss for keeping the alignment between instruction and image. Let H^{V+} and H^{V-} represent the positive and negative examples, the loss \mathcal{L}^{Coh} of the i -th step can be defined as follows:

$$\mathcal{L}_i^{\text{Coh}} = -\log \frac{\exp\{\text{sim}(\mathbf{H}_i^T, \mathbf{H}_i^{V+})/\tau\}}{\sum_{j=1}^K \exp\{\text{sim}(\mathbf{H}_i^T, \mathbf{H}_j^{V-})/\tau\}}, \quad (4)$$

where K is the total number of negative samples (He et al., 2020) generated from the min-batch, $\text{sim}(\cdot, \cdot)$ and τ are the standard cosine similarity function and temperature.

In a downstream reasoning task, the model needs to predict the correct candidate $a_j = \{I_{j,1}, \dots, I_{j,N_a}\}$ based on the instructions $\mathcal{S} = \{s_t\}_{t=1}^{N_t}$. Referring to the sentence image predic-

tion task in (Li et al., 2020), we concatenate all representations of each candidate image sequence to generate a instruction candidate pair as: $(\mathbf{S}, a_j) = [\text{CLS}, H_1^T, \dots, H_{N_t}^T, \text{SEP}, H_{j,1}^V, \dots, H_{j,N_a}^V]$, where [CLS] and [SEP] are special tokens as used in (Li et al., 2020). We pass this input through a shallow transformer followed by a fully connected layer to obtain the prediction score $P(\mathbf{S}, a_j)$ for the j -th candidate, and the prediction loss can be defined as

$$\mathbf{L}^{\text{Pre}} = -\log \frac{\exp(P(\mathbf{S}, a_j))}{\sum_{i=1, i \neq j}^{N_a-1} \exp(P(\mathbf{S}, a_i))}, \quad (5)$$

where a_j is the correct candidate and N_a is the number of candidates. We get the final loss function and optimize it through the Adam optimizer:

$$\mathbf{L} = \mathbf{L}^{\text{Pre}} + \lambda_b \mathbf{L}^{\text{Coh}}, \quad (6)$$

where λ_b is the balance parameter. Unless otherwise specified, all the results in this paper use $\lambda_b = 0.1$ which we find to perform best.

4 Experiments

4.1 Datasets and Metrics

RecipeQA. RecipeQA (Yagcioglu et al., 2018) is a multimodal comprehension dataset with 20K recipes approximately and more than 36K question-answer pairs. Unlike other multimodal reading comprehension datasets (Tapaswi et al., 2016; Iyyer et al., 2017; Kembhavi et al., 2017) analyze against movie clips or comics, RecipeQA requires reasoning real-world cases.

CraftQA. We collect CraftQA from Instructables², which is an online community where people can share their tutorials for accomplishing a task in a step-by-step manner. Specifically, we collect the most visited tutorials and remove those that contain only text or video. For question and answer generation, we also remove the tutorials that contain less than 3 images. To construct the distractor of each task, we compute the Euclidean distance between the image features that are extracted from a pretrained ResNet-50 (He et al., 2016). Taking the visual cloze task as an example, the distractor is sampled from the nearest neighbors of the ground-truth image based on Euclidean distance. Finally, CraftQA contains about 27k craft product-making tutorials and 46k question-answer pairs. We employ CraftQA to evaluate the reading comprehension performance of TMEG in different domains as

²<https://www.instructables.com/>

Dataset	Statistics	Train	Valid	Test
RecipeQA	# of recipes	15,847	1,963	1,969
	avg. # of steps	5.99	6.01	6.00
	avg. # of words	443.01	440.51	435.33
	avg. # of images	12.67	12.74	12.65
CraftQA	# of tutorials	21,790	2,601	2,604
	avg. # of steps	7.53	7.47	7.56
	avg. # of words	535.88	531.01	541.97
	avg. # of images	20.14	20.19	20.37

Table 1: Statistics of RecipeQA and CraftQA dataset. Each dataset is split into training, validation, and test sets based on the number of recipes or craft-making tutorials. We also provide the average count of steps, images, and words contained in each split dataset.

well as its domain transfer capability. More statistics about these two datasets are shown in Table 1.

Metric. In three Procedural M³C tasks that are tested in the following experiments (visual cloze, visual coherence, and visual ordering), we use classification accuracy as the evaluation metric, which is defined as the percentage of yielding the ground-truth answer during testing (Yagcioglu et al., 2018; Amac et al., 2019; Liu et al., 2020).

4.2 Implementation Details

For visual node construction, we employ the pre-trained Faster R-CNN (Ren et al., 2015) model provided by Detectron2 (Wu et al., 2019) and limit the number of objects to 36 for each image. Following (Yang et al., 2019; Song et al., 2021), we set the thresholds λ_t and λ_m as 7 and 0.5, respectively, for the temporal and the modal edge constructions.

The framework of the graph-based fusion module is built on VisualBERT (Li et al., 2020) with its initialized parameters and tokenizer implemented by HuggingFace’s transformers library (Wolf et al., 2020). The shallow transformer in the graph-based reasoning module is designed as 2 hidden layers with a size of 512 and 8 attention heads. During the training stage, the batch size is fixed to 16 and the number of negative samples K is set to 8. The temperature parameter τ in Eq.(4) is set to 0.07. The balance parameter λ_b in Eq.(6) is set to 0.1. Adam with the learning rate 5×10^{-5} is used to update parameters. We introduce an early stopping mechanism and set the patience value to 5, which means the training will stop if the model performance is not improved in five consecutive times. Our source code will be released online.

Model	RecipeQA				CraftQA			
	Cloze	Coherence	Ordering	Average	Cloze	Coherence	Ordering	Average
Human (Amac et al., 2019)	77.60	81.60	64.00	74.40	55.33	61.80	54.60	57.24
HS (Yagcioglu et al., 2018)	27.35	65.80	40.88	44.68	27.81	42.61	35.23	35.22
PRN (Amac et al., 2019)	56.31	53.64	62.77	57.57	35.14	33.67	42.31	37.04
MLMM-Trans (Liu et al., 2020)	65.57	67.33	63.75	65.55	39.43	46.52	43.62	43.20
VisualBERT (Li et al., 2020)	66.49	68.15	63.09	65.91	40.56	47.84	46.18	44.86
TMEG (Our Model)	73.27	70.38	65.56	69.73	48.01	51.36	49.25	49.54
w/o Temporal Encoding	70.18	68.81	63.91	67.63	45.91	47.71	47.24	46.95
w/o Modal Encoding	71.50	69.63	64.15	68.42	46.67	49.06	48.35	48.02
w/o Both Encoding	68.97	67.72	63.34	66.67	44.87	46.29	46.77	45.98
w/o Contrastive Loss \mathcal{L}^{Coh}	71.66	68.54	64.54	68.24	47.49	50.03	48.79	48.77

Table 2: Experimental comparison of procedural multimodal machine comprehension on RecipeQA and CraftQA: “w/o Temporal Encoding” and “w/o Modal Encoding” denote to remove the temporal or modal edges respectively and “w/o Both Encoding” denotes to remove both of them. “w/o Contrastive Loss \mathcal{L}^{Coh} ” denotes excluding the contrastive coherence loss \mathcal{L}^{Coh} from the final loss. Similar to RecipeQA, 100 questions in CraftQA are extracted from its validation set to evaluate the “Human” performance.

4.3 Baselines

We compare our model with the following models: (1) **Hasty Student (HS)** (Yagcioglu et al., 2018) discards textual context and directly exploits the similarities and dissimilarities between answer images to rank candidates. (2) **PRN** (Amac et al., 2019) introduces external relational memory units to keep track of textual entities and employs a bi-directional attention mechanism to obtain a question-aware embedding for prediction. (3) **MLMM-Trans** (Liu et al., 2020) modifies the framework of the transformer (Vaswani et al., 2017) and conducts an intensive attention mechanism at multiple levels to predict correct image sequences. (4) **VisualBERT** (Li et al., 2020) consists of a stack of transformer layers that extend the traditional BERT (Devlin et al., 2019) model to a multimodal encoder. The performance of some baselines on RecipeQA has been previously reported in (Amac et al., 2019; Liu et al., 2020).

4.4 Experimental Results

4.4.1 Comparison Analysis

As shown in Table 2, TMEG shows favorable performance in different reasoning tasks, with an average accuracy of 69.73 and 49.54 on RecipeQA and CraftQA, following behind the “Human” performance. Besides, the performance on the *visual ordering* task exceeds human accuracy for the first time, which proves that the temporal and modal analysis in TMEG is effective in comprehending PMDs. MLMM-Trans performs comparably with VisualBERT while inferior to TMEG, which may be attributed to their superficial consideration of




Figure 3: Experimental results of our model with different balance parameter λ_b in Eq.(6) on RecipeQA.

entities.

MLMM-Trans ignores the entity information contained in text (e.g., the correspondence between entities in text and images) and VisualBERT directly fuses textual and visual features without considering entity evolution. In TMEG, we explicitly identify and model entity evolution in PMD, whereas MLMM-Trans and VisualBERT assume entity information to be learned implicitly alongside other data.

Meanwhile, CraftQA has more images (20.14 vs 12.67) and tokens (535.88 vs 443.01) on average than RecipeQA. More diverse complex cases in CraftQA require better comprehension and reasoning capacities for both models and humans. We believe this can explain the lower results on CraftQA. This emphasizes the necessity of comprehending entity coherence in a multimodal context.

4.4.2 Ablation Study

We evaluate the effects of temporal encoding, modal encoding, and contrastive coherence loss

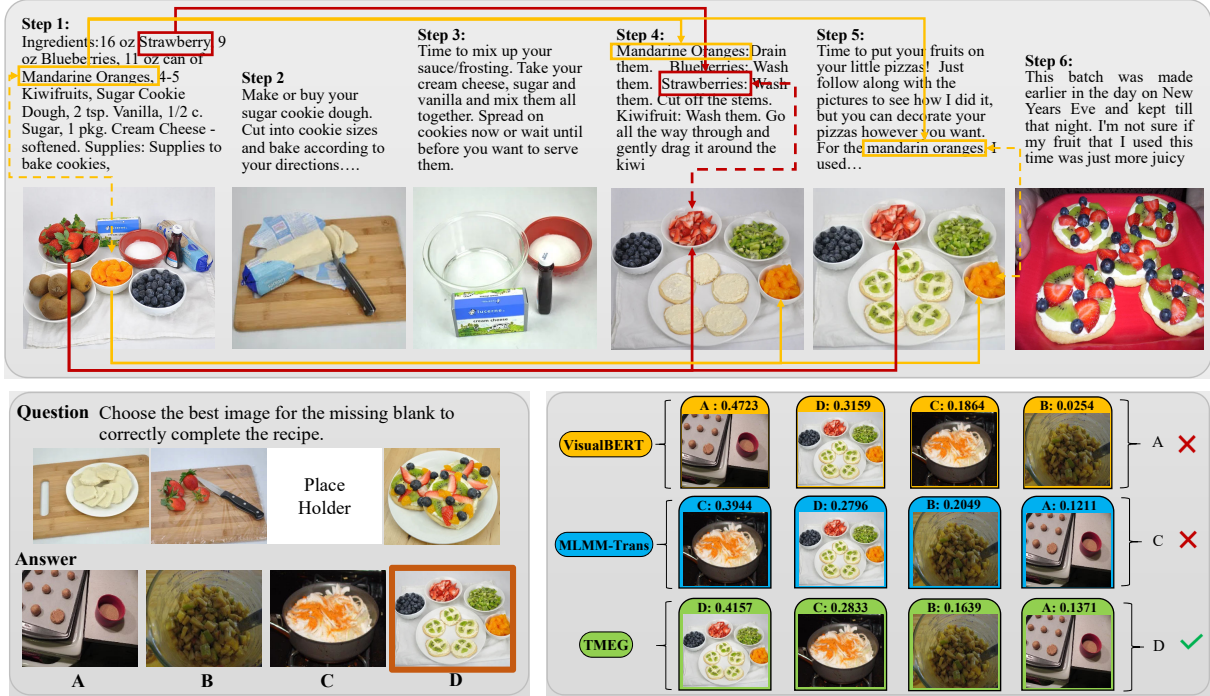


Figure 4: An illustrated example of the visual cloze task on RecipeQA to clarify the workflow of TMEG when dealing with the downstream tasks.

\mathcal{L}^{Coh} to examine the three modules.

Edge Encoding. Table 2 also shows the ablation results of our model when each module is respectively removed. In terms of edge encoding, removing temporal encoding makes more negative effects on TMEG than moving the modal encoding, reflecting the significance of modeling temporal entity evolution for procedural M^3C .

Contrastive Coherence Loss. As shown in the last row of Table 2, we find that \mathcal{L}^{Coh} can indeed improve TMEG. The reason is that \mathcal{L}^{Coh} helps enhance the learning of textual entity representation and visual entity representation in Procedural M^3C .

4.5 Analysis of TMEG

4.5.1 Balance Parameter

In Figure 3, we illustrate the influence of the balance parameter λ_b in Eq.(6), which balances the contrastive coherence loss \mathcal{L}^{Coh} and the candidate prediction loss \mathcal{L}^{Pre} . We tune λ_b from 0 to 0.2 with 0.05 as the step size. We observe that the model beats the highest accuracy when $\lambda_b = 0.1$. Generally, (1) introducing the contrastive coherence loss can improve TMEG for better fitting downstream tasks, and (2) appropriately balancing the prediction loss \mathcal{L}^{Pre} and contrastive coherence loss \mathcal{L}^{Coh} helps TMEG comprehend PMDs.

Model	R2C	C2R	Average
MLMM-Trans (Liu et al., 2020)	33.98	40.14	37.06
VisualBERT (Li et al., 2020)	35.24	42.15	38.69
TMEG (Our Model)	39.06	45.42	42.24
TMEG (w/o Edge Encoding)	36.02	43.17	39.60

Table 3: Results of the domain transfer experiments, where ‘‘R2C’’ denotes training models on RecipeQA while testing on CraftQA, and ‘‘C2R’’ represents the opposite.

4.5.2 Cross-Domain Investigation

To study the domain transfer capability of our framework, we evaluate TMEG in different domains, as shown in Table 3. Specifically, The model trained on RecipeQA is evaluated on CraftQA, and the reverse is true for CraftQA. Results show that compared with other baselines, our model achieves more generalized and better comprehension performance on domain transfer by incorporating TMEG.

4.5.3 Case Study

Figure 4 further presents a visual cloze example on RecipeQA which requires a correct image in the missing piece after reading the context. We compare the highest-scored candidate images respectively picked out by MLMM-Trans (Liu et al., 2020), VisualBERT (Li et al., 2020), and TMEG.

By considering the temporal-modal entity evolution, TMEG can capture the salient entities (e.g., Strawberry and Sugar Cookie Dough) and trace their evolution at each step, thereby inferring the ground-truth answer.

5 Conclusion

In this paper, we propose a novel temporal-modal entity graph (TMEG) to approach Procedural M³C. Based on TMEG, we introduce graph-based fusion module and reasoning module, which are used to aggregate node features and solve downstream reasoning tasks.

What’s more, we introduce another Procedural M³C dataset called CraftQA to assist in evaluating the generalization performance of TMEG in different domains and domain transfer. Extensive experiments on the RecipeQA and CraftQA validate the superiority of TMEG. A promising future direction is to introduce temporal-modal entity graphs into the video understanding task (Lin et al., 2020; Xu et al., 2020), which also calls for an enhancement of the temporal and the cross-modal reasoning capability.

Ethical Considerations

Intellectual Property. CraftQA contains question answer pairs generated from copyright free tutorials found online³. All of the tutorials are licensed with the Creative Commons license⁴ which helps share knowledge and creativity for common use.

The collection of CraftQA is in accordance with the Terms of Service of Instructables as follows: by posting, providing, uploading, submitting, sharing, publishing, distributing, making available or allowing others to access and/or use Your Content to or through the Service You are solely responsible and liable for the consequences of doing so and you acknowledge and agree that Your Content can and may be viewed worldwide⁵.

We also construct experimental evaluations on the RecipeQA dataset. Referring to the official dataset descriptions of RecipeQA⁶, Legal and Ethical Considerations had been taken into account during the construction of RecipeQA. We have cited the corresponding papers in this study.

³<https://www.instructables.com>

⁴<https://creativecommons.org/licenses/by-nc-sa/4.0>

⁵<https://www.autodesk.com/company/legal-notices-trademarks/terms-of-service-autodesk360-web-services/instructables-terms-of-service-june-5-2013>

⁶<https://hucv1.github.io/recipeqa/recipeqa-datasheet.pdf>

Privacy. According to the Privacy Statement of Instructables⁷, users can choose whether or not to expose their information when publishing tutorials. Respecting personal privacy, we have removed all of the personal information of users from CraftQA and promise CraftQA isn’t involved with any privacy issues.

Acknowledgements: This work was supported in part by the National Natural Science Foundation of China (No. 62106091) and Shandong Provincial Natural Science Foundation (No. ZR2021MF054).

References

- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *Proceedings of the 27th International Conference on Computational, pages* 1638–1649.
- Malihe Alikhani, Fangda Han, Hareesh Ravi, Mubbasir Kapadia, Vladimir Pavlovic, and Matthew Stone. 2021. Cross-modal coherence for text-to-image retrieval. *CoRR*, abs/2109.11047.
- Mustafa Sercan Amac, Semih Yagcioglu, Aykut Erdem, and Erkut Erdem. 2019. Procedural reasoning networks for understanding multimodal procedures. In *Proceedings of the 23rd Conference on Computational Natural Language Learning*, pages 441–451.
- Aida Amini, Antoine Bosselut, Bhavana Dalvi Mishra, Yejin Choi, and Hannaneh Hajishirzi. 2020. Procedural reading comprehension with attribute-aware context flow. In *Proceedings of the Conference on Automated Knowledge Base Construction*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2425–2433.
- Antoine Bosselut, Omer Levy, Ari Holtzman, Corin Ennis, Dieter Fox, and Yejin Choi. 2018. Simulating action dynamics with neural process networks. In *Proceedings of the 6th International Conference on Learning Representations*.
- Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. Connecting the dots: Document-level neural relation extraction with edge-oriented graphs. *arXiv preprint arXiv:1909.00228*.
- Rajarshi Das, Tsendsuren Munkhdalai, Xingdi Yuan, Adam Trischler, and Andrew McCallum. 2019. Building dynamic knowledge graphs from text using machine reading comprehension. In *Proceedings of the 7th International Conference on Learning Representations*.
- ⁷<https://www.autodesk.com/company/legal-notices-trademarks/privacy-statement>

- 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 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the 9th International Conference on Learning Representations*.
- Hossein Rajaby Faghihi and Parisa Kordjamshidi. 2021. Time-stamped language model: Teaching language models to understand the flow of events. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 4560–4570. Association for Computational Linguistics.
- Shaohua Fan, Junxiong Zhu, Xiaotian Han, Chuan Shi, Linmei Hu, Biyu Ma, and Yongliang Li. 2019. Metapath-guided heterogeneous graph neural network for intent recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2478–2486.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9726–9735.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.
- Jack Hessel, Lillian Lee, and David Mimno. 2019. Unsupervised discovery of multimodal links in multi-image, multi-sentence documents. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2034–2045.
- Hao Huang, Xiubo Geng, Jian Pei, Guodong Long, and Daxin Jiang. 2021. Reasoning over entity-action-location graph for procedural text understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and 11th International Joint Conference on Natural Language Processing*, pages 5100–5109.
- Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan L. Boyd-Graber, Hal Daumé III, and Larry S. Davis. 2017. The amazing mysteries of the gutter: Drawing inferences between panels in comic book narratives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6478–6487.
- Aniruddha Kembhavi, Min Joon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5376–5384.
- Mahmoud Khademi. 2020. Multimodal neural graph memory networks for visual question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7177–7188.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2020. What does BERT with vision look at? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5265–5275.
- Angela S. Lin, Sudha Rao, Asli Celikyilmaz, Elnaz Nouri, Chris Brockett, Debadeepta Dey, and Bill Dolan. 2020. A recipe for creating multimodal aligned datasets for sequential tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4871–4884.
- Ao Liu, Shuai Yuan, Chenbin Zhang, Congjian Luo, Yaqing Liao, Kun Bai, and Zenglin Xu. 2020. Multi-level multimodal transformer network for multimodal recipe comprehension. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 1781–1784.
- Nicola Messina, Fabrizio Falchi, Andrea Esuli, and Giuseppe Amato. 2020. Transformer reasoning network for image-text matching and retrieval. In *Proceedings of the 25th International Conference on Pattern Recognition*, pages 5222–5229.
- Bhavana Dalvi Mishra, Niket Tandon, Antoine Bosselut, Wen-tau Yih, and Peter Clark. 2019. Everything happens for a reason: Discovering the purpose of actions in procedural text. *CoRR*, abs/1909.04745.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015*, pages 91–99.
- Xue Song, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. 2021. Spatial-temporal graphs for cross-modal text2video retrieval. *IEEE Transactions on Multimedia*.
- Niket Tandon, Keisuke Sakaguchi, Bhavana Dalvi, Dheeraj Rajagopal, Peter Clark, Michal Guerquin, Kyle Richardson, and Eduard H. Hovy. 2020. A

- dataset for tracking entities in open domain procedural text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 6408–6417.
- Makarand Tapaswi, Yukun Zhu, Rainer Stiefelwagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4631–4640.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pages 5998–6008.
- Jesse Vig. 2019. A multiscale visualization of attention in the transformer model. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 37–42.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45.
- Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Frank F. Xu, Lei Ji, Botian Shi, Junyi Du, Graham Neubig, Yonatan Bisk, and Nan Duan. 2020. A benchmark for structured procedural knowledge extraction from cooking videos. *CoRR*, abs/2005.00706.
- Semih Yagcioglu, Aykut Erdem, Erkut Erdem, and Nazli Ikizler-Cinbis. 2018. Recipeqa: A challenge dataset for multimodal comprehension of cooking recipes. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1358–1368.
- Jianing Yang, Yongxin Wang, Ruitao Yi, Yuying Zhu, Azaan Rehman, Amir Zadeh, Soujanya Poria, and Louis-Philippe Morency. 2021. MTAG: modal-temporal attention graph for unaligned human multimodal language sequences. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1009–1021.
- Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. 2019. A fast and accurate one-stage approach to visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4682–4692.
- Yongjing Yin, Fandong Meng, Jinsong Su, Chulun Zhou, Zhengyuan Yang, Jie Zhou, and Jiebo Luo. 2020. A novel graph-based multi-modal fusion encoder for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3035.
- Dong Zhang, Suzhong Wei, Shoushan Li, Hanqian Wu, Qiaoming Zhu, and Guodong Zhou. 2021a. Multi-modal graph fusion for named entity recognition with targeted visual guidance. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence and 33rd Conference on Innovative Applications of Artificial Intelligence and 7th Symposium on Educational Advances in Artificial Intelligence*, pages 14347–14355.
- Zhihan Zhang, Xiubo Geng, Tao Qin, Yunfang Wu, and Daxin Jiang. 2021b. Knowledge-aware procedural text understanding with multi-stage training. In *Proceedings of the International World Wide Web Conference 2021*, pages 3512–3523.
- Wanjun Zhong, Duyu Tang, Nan Duan, Ming Zhou, Jiahai Wang, and Jian Yin. 2020. A heterogeneous graph with factual, temporal and logical knowledge for question answering over dynamic contexts. *CoRR*, abs/2004.12057.