
Building Dynamic Knowledge Graphs from Text-based Games

Mikuláš Zelinka^{†,*}, Xingdi Yuan^{‡,*}, Marc-Alexandre Côté^{‡,*}

Romain Laroche[‡], Adam Trischler[‡]

[†]Charles University, Faculty of Mathematics and Physics, Czech Republic

[‡]Microsoft Research, Montréal

eric.yuan@microsoft.com

Abstract

We are interested in learning how to update Knowledge Graphs (KG) from text. In this preliminary work, we propose a novel Sequence-to-Sequence (Seq2Seq) architecture to generate elementary KG operations. Furthermore, we introduce a new dataset for KG extraction built upon text-based game transitions (over 300k data points). We conduct experiments and discuss the results.

1 Introduction

Text-based games are complex, interactive simulations in which text describes the game state and players make progress by entering text actions. They can be seen as sequential decision making tasks where accomplishing certain goals earns rewards (points). Solving these games requires both Reinforcement Learning (RL) and Natural Language Processing (NLP) techniques.

Given the complex, partially observable nature of text-based games, an explicit structured memory – e.g., in the form of a graph – is a useful component for game-playing agents. In this work, we side step the game-playing aspect of the problem to focus solely on learning how to build and dynamically maintain a KG from text observations. Specifically, our proposed model learns to generate graph update operations to update an existing KG given new text information. We see this KG module as an independent block that can be leveraged by game-playing agents to improve their performance.

Related Work: Numerous recent works focus on learning or using KGs in textual environments. Das et al. (2018) leverage a machine reading comprehension (MRC) mechanism to query for entities and states in short text passages and use attention to address aliased entity occurrences and to track the entity states dynamically. Xiong et al. (2018) focus on one-shot learning of new relations from only one training instance.

For text games specifically, Ammanabrolu and Riedl (2018) leverage KGs to improve performance, relying on OpenIE for entity extraction and several game-specific rules for building and maintaining the KG. Sinha et al. (2019) introduce a dataset aimed at natural language understanding and generalization in reasoning about entity relations that was built using a KG-like structure.

To the best of our knowledge, all of the approaches for learning KGs are either concerned with building static KGs (rather than focusing on small, dynamic updates), or employ some domain-specific knowledge or rules to facilitate learning. In contrast, our proposed model learns to generate general update operations for modifying a KG.

*Equal contribution, work performed when first author was an intern at Microsoft Research, Montréal.




Figure 1: Illustration of an example in TEXTWORLD KG. By issuing an action A_{t-1} at game step $t - 1$, the environment returns a new observation, O_t . Given the KG at step $t - 1$, a model is required to predict the new KG given the text observation.

2 The TextWorld KG Dataset

In this section, we introduce a new dynamic KG extraction dataset, TEXTWORLD KG. TEXTWORLD KG is based on a set of text-based games generated using TextWorld (Côté et al., 2018).² That framework enables us to extract the underlying *partial* KG for every state, i.e., the subgraph that represents the agent’s partial knowledge of the world – what it has observed so far. All games share the same overarching theme: the agent finds itself hungry in a simple modern house with the goal of gathering ingredients and cooking a meal.

To build the TEXTWORLD KG dataset, we collect game transitions obtained by following each game’s walkthrough (provided by TextWorld). Additionally, after each step in a walkthrough, we perform 5 additional actions sampled at random from the list of admissible commands³ (also provided by TextWorld). This presumably promotes robustness and generalizability of a training agent since it will encounter off-the-path transitions during game playing in the RL setting due to the absent of walkthroughs. Therefore an agent pre-trained on such data is more likely to work well in the RL setting. Formally, each data point in TEXTWORLD KG is a tuple, $\{G_{t-1}^{\text{seen}}, A_{t-1}, O_t, G_t^{\text{seen}}\}$, where G_{t-1}^{seen} is a partial KG (in the format of Resource Description Framework (RDF) triples) representing the information an agent has seen during all previous game steps up to $t - 1$. After the agent issues an action A_{t-1} (a string of words), the game engine returns a new text observation O_t describing its effects on the world. The task is to predict the updated partial KG G_t^{seen} given the above information. Note that the new observation might not contain new information, since some actions do not change the game state (e.g. look in the same room twice). Table 1 shows some statistics about TEXTWORLD KG and Figure 1 illustrates an example data point.

An important challenge posed by TextWorld is generalization. In each individual game instance, the interactable objects and their locations change along with the layout of the environment. Similarly, object names can be composed of multiple adjectives and a noun (e.g., red hot chili pepper), and at test time, players may encounter object names never seen during training. TEXTWORLD KG inherits both of these challenging features.

²We use the dataset provided for the *First TextWorld Problems* competition, available at <http://aka.ms/ftwp-dataset>.

³This is the set of game actions understood by the game in a given state.

#Train	#Valid	#Test	Avg. Obs.	Avg. #Operations	#Vertices	#Edges	Avg. #Connections
267,031	13,442	41,865	29.3 tokens	3.1	99	10	43.1

Table 1: Statistics of TEXTWORLD KG. **Avg. Obs.** is the average number of tokens an observation has. **Avg. #Operations** is the average number of update operations to generate per time step. **#Vertices** and **#Edges** correspond to the number of unique entities and relation types. **Avg. #Connections** is the average number of connections a graph has.

3 Learning to Update a KG

3.1 KG Definition

In a text-based game, at any given game step t , the game state s_t can be represented as a graph $\mathcal{G}_t^{\text{full}} = (\mathcal{V}_t, \mathcal{E}_t)$. In our setting, vertices \mathcal{V}_t represent entities (including objects, the player, and locations) and their states (e.g., closed, fried, sliced). Vertices are connected by edges \mathcal{E} , which represent a set of relations between entities (e.g. north_of, in, is).

Since games are partially observable, at every step an agent only observes part of the full game state (e.g., the agent cannot know facts in a room it has not visited). Thus, an agent must build its belief about the world, $\mathcal{G}_t^{\text{belief}}$, from its observations. Ideally, the belief graph should match the ground truth graph, $\mathcal{G}_t^{\text{seen}}$, which is a subgraph of $\mathcal{G}_t^{\text{full}}$ representing what has been seen so far in the game.

TextWorld games are deterministic, so by progressively exploring and observing, an agent should discover more knowledge to push into its belief graph. Eventually, this ought to converge to a graph that accurately represents the entire game state.

3.2 Updating a KG

Instead of generating the entire belief graph at every game step, we generate a set of update operations Δg_t such that $\mathcal{G}_t^{\text{belief}} = \text{Update}(\mathcal{G}_{t-1}^{\text{belief}}, \Delta g_t)$, where Update is an oracle function that applies Δg_t . In our case, each update operation in Δg_t is represented as a text command. We define the following two types of update operation:

- `add(node1, node2, relation)`: add a directed edge, named `relation`, between `node1` and `node2`; if any of these nodes does not exist, add that node first.
- `delete(node1, node2, relation)`: delete a directed edge, named `relation`, between `node1` and `node2`; if any of the nodes or the edge does not exist, ignore this operation.

Given a new observation string O_t and an agent’s current belief $\mathcal{G}_{t-1}^{\text{belief}}$, the agent is required to generate $k \geq 0$ operations as defined above to merge newly observed information into its belief graph. For the example shown in Figure 1, the generated operations are listed in Table 2.

We formulate the update generation task as a Seq2Seq problem. Specifically, we adopt the decoding strategy from Yuan et al. (2018), where given an observation sequence O_t and a belief graph $\mathcal{G}_{t-1}^{\text{belief}}$, the agent generates a sequence of tokens consisting of multiple graph update operations separated by a delimiter token.

As pointed out by Meng et al. (2019), the order of ground-truth tokens and sequences (in our case, graph update operations) matters in Seq2Seq language generation. We therefore define a set of rules (e.g., always add before delete) to order ground-truth operations for teacher forcing during training.

```

add (player, shed, at)
add (shed, backyard, west_of)
add (wooden door, shed, east_of)
add (toolbox, shed, in)
add (toolbox, closed, is)
add (workbench, shed, in)
delete (player, backyard, at)

```

Table 2: Update operations corresponding to the transition shown in Figure 1.

3.3 Model Architecture

We use a transformer-based Seq2Seq model (Vaswani et al., 2017) to generate update operations. As shown in Figure 2 in Appendix A, the model consists of the following components:

- A **text encoder**, which reads text inputs (the concatenation of observation O_t and the action at the

previous game step, A_{t-1}), and generates hidden representations.

- A **graph encoder**, which encodes the previous belief $\mathcal{G}_{t-1}^{\text{belief}}$ into hidden representations.
- An attention-based **representation aggregator**, which combines the two above representations.
- A **command generator**, which takes aggregated representations and generates update operations token by token.

For space considerations, we elaborate our model components in Appendix A. Following common practice in natural language generation (NLG), we train our operation generation model via teacher forcing. Specifically, during training, a right-shifted ground truth target sequence is provided as input to the decoder and the model is trained with the negative log-likelihood (NLL) loss. During test, the model starts generating from a start-of-sentence token and uses the previously generated token as input to the next step. The model terminates after generating an end-of-sequence token.

4 Experiments and Discussion

In this preliminary study we test 4 graph encoder variants of the proposed model. First, as a baseline, we disable the graph encoder, which renders the model a standard Seq2Seq transformer. Second, we utilize a Graph Convolutional Network (GCN) (Kipf and Welling, 2016) as the graph encoder. The GCN does not consider multiple relations.⁴ Third, we enable conditioning on multiple relations by using a Relational Graph Convolutional Network (R-GCN) (Schlichtkrull et al., 2018). Although R-GCN takes into account multiple relations, it does not consider information in relation labels. In our task, this information is important (e.g., `east_of` and `west_of` are symmetric relations). Therefore, we finally learn a vector representation for each relation that is conditioned on the label’s text embeddings. The resulting relation representation is used as an extra input to the R-GCN layer. Table 3 shows the test results for all models.

Model	TF-F ₁	FR-F ₁
Transformer	0.832	0.434
+ GCN	0.928	0.664
+ R-GCN	0.965	0.645
+ R-Emb	0.962	0.697

Table 3: Test performance.

During training, a model takes $\{\mathcal{G}_{t-1}^{\text{seen}}, A_{t-1}, O_t\}$ as input, where $\mathcal{G}_{t-1}^{\text{seen}}$ is the input graph, and A_{t-1} and O_t are the text action issued at the previous game step and the resulting text observation, respectively. The model outputs a sequence describing an update operation to the graph and the resulting $\mathcal{G}_t^{\text{belief}}$. During evaluation:

- **Teacher-force (TF) F₁**: we use the ground-truth $\mathcal{G}_{t-1}^{\text{seen}}$ as input graph and compute the F₁ score between the model’s generated graph update *commands* and the ground-truth commands. Note the F₁ score is computed on command level (i.e., if any token in a command is incorrect, this command is treated as incorrect).
- **Free-run (FR) F₁**: we initialize the belief graph at the beginning of each game with an empty graph. For each game step, we use $\mathcal{G}_{t-1}^{\text{belief}}$ (the graph generated by the model) as input. At the end of each game, we compute F₁ score between the final belief graph $\mathcal{G}_T^{\text{belief}}$ and ground truth $\mathcal{G}_T^{\text{seen}}$, graphs are represented as *RDF triples*.

In general, although all model variants show good performance on **TF-F₁**, they perform worse on **FR-F₁**. This is not surprising since errors accumulate in the latter setting. Models using R-GCN outperform those using GCN by a noticeable margin, which suggests relational information is essential in the proposed tasks. Interestingly, while the two R-GCN models perform similarly on **TF-F₁**, the variant with relational embedding (considering information in relation labels) significantly outperforms the other on **FR-F₁**.

To better understand the behavior of our proposed models on `TEXTWORLD KG`, we conduct an error analysis, which we show in Appendix B.

The next step for this project is to leverage the KG update module while playing text-based games. We believe that maintaining such a graph could help an RL agent (1) to avoid re-discovering known facts about the world and (2) to discover new world knowledge efficiently. We are also interested in finding ways of transferring learned graphs from one game to another to improve agents’ ability to generalize.

⁴For models that do not consider relational information, we use single relation KGs as ground-truth during evaluation.

References

- Ammanabrolu, P. and Riedl, M. O. (2018). Playing text-adventure games with graph-based deep reinforcement learning. *CoRR*, abs/1812.01628.
- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Côté, M.-A., Kádár, A., Yuan, X., Kybartas, B., Barnes, T., Fine, E., Moore, J., Hausknecht, M., Asri, L. E., Adada, M., Tay, W., and Trischler, A. (2018). Textworld: A learning environment for text-based games.
- Das, R., Munkhdalai, T., Yuan, X., Trischler, A., and McCallum, A. (2018). Building dynamic knowledge graphs from text using machine reading comprehension. *CoRR*, abs/1810.05682.
- Gülçehre, Ç., Ahn, S., Nallapati, R., Zhou, B., and Bengio, Y. (2016). Pointing the unknown words. *CoRR*, abs/1603.08148.
- Kipf, T. N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907.
- Meng, R., Yuan, X., Wang, T., Brusilovsky, P., Trischler, A., and He, D. (2019). Does order matter? an empirical study on generating multiple keyphrases as a sequence. *CoRR*.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., Van Den Berg, R., Titov, I., and Welling, M. (2018). Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.
- Seo, M., Kembhavi, A., Farhadi, A., and Hajishirzi, H. (2016). Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Sinha, K., Sodhani, S., Dong, J., Pineau, J., and Hamilton, W. L. (2019). Clutrr: A diagnostic benchmark for inductive reasoning from text. *arXiv preprint arXiv:1908.06177*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. *CoRR*, abs/1706.03762.
- Xiong, W., Yu, M., Chang, S., Guo, X., and Wang, W. Y. (2018). One-shot relational learning for knowledge graphs. *arXiv preprint arXiv:1808.09040*.
- Yuan, X., Wang, T., Meng, R., Thaker, K., Brusilovsky, P., He, D., and Trischler, A. (2018). Generating diverse numbers of diverse keyphrases. *CoRR*, abs/1810.05241.

A Model Architecture




Figure 2: Graph update operation generation model.

Text Encoder: An observation string O_t is provided in response to the text action A_{t-1} an agent issued at previous game step. We concatenate the two text information together $[A_{t-1}; O_t]$ as the input to the text encoder, in which $[\cdot; \cdot]$ indicates vector/string concatenation. The text encoder consists an embedding layer and a stack of transformer blocks. The text encoder results a sequence of hidden vectors $h^O \in \mathbb{R}^{L^O \times H}$, where L^O is length of the concatenated string, H is hidden size.

Graph Encoder: At the same time, the graph encoder takes the agent’s belief KG $\mathcal{G}_{t-1}^{\text{belief}}$ (in which stores the agent’s observations at all previous game steps) as input. We adopt different off-the-shelf graph neural networks (GNN) as the graph encoder (we will describe more details in experiment section). After several layers of propagation, graph representations $h^{\mathcal{G}} \in \mathbb{R}^{N \times H}$ are generated, where N is number of nodes in the KG.

Representation Aggregator: We use an attention based layer to aggregate text and graph representations (Bahdanau et al., 2014; Seo et al., 2016). Specifically, we use weighted sum of text representations to represent graph information, which results $h^{\mathcal{G}O} \in \mathbb{R}^{N \times H}$; similarly, we use weighted sum of graph representations to represent text information, resulting $h^{OG} \in \mathbb{R}^{L^O \times H}$.

Command Generator Finally, both h^{OG} and $h^{\mathcal{G}O}$ are used to condition text generation. Input tokens are first converted into embeddings, then they are fed into a stack of decoder transformer blocks which result a probability distribution over the vocabulary. To prevent model from utilizing future information, we follow Vaswani et al. (2017) to use a masked multi-head self attention layer in the beginning of each block. For being able to both generate a word from vocabulary, and point a word from the source text O_t , we adopt the pointer-softmax mechanism (Gülçehre et al., 2016).

B Error Analysis




Figure 3: Average **TF-F₁** scores grouped by verbs.

In Figure 3, we report average **TF-F₁** scores grouped by verbs found in input actions A_{t-1} . We can observe that vanilla transformer model performs poorly on `go` actions, which aligns with the fact that after issuing a `go` action, the resulting observation text O_t does not contain any information of the agent’s previous location. On the other hand, the other models benefit from their belief graph to retrieve that *single* information.

Also, we notice that all models perform relatively poorly on `prepare` actions. This also makes sense since in TextWorld games, the action of preparing a meal consumes multiple food ingredients at once in order to produce a meal object. The resulting observation text O_t following a `prepare` action only contains information about the newly produced meal and does not mention what food ingredients have been consumed. Even though the information about the recipe (i.e. ingredients needed) is part of the KG graph, a model has to learn to retrieve *multiple* information at once from its belief graph $\mathcal{G}_{t-1}^{\text{belief}}$ to figure out what ingredients will be consumed.