

---

# Compositional Language Understanding with Text-based Relational Reasoning

---

**Koustuv Sinha\***

Mila

McGill University, Canada

koustuv.sinha@mail.mcgill.ca

**Shagun Sodhani**

Mila

Université de Montréal, Canada

sshagunsodhani@gmail.com

**William L. Hamilton**

Mila

McGill University, Canada

Facebook AI Research (FAIR), Montreal

wlh@cs.mcgill.ca

**Joelle Pineau**

Mila

McGill University, Canada

Facebook AI Research (FAIR), Montreal

jpineau@cs.mcgill.ca

## Abstract

Neural networks for natural language reasoning have largely focused on extractive, fact-based question-answering (QA) and common-sense inference. However, it is also crucial to understand the extent to which neural networks can perform relational reasoning and combinatorial generalization from natural language—abilities that are often obscured by annotation artifacts and the dominance of language modeling in standard QA benchmarks. In this work, we present a novel benchmark dataset for language understanding that isolates performance on relational reasoning. We also present a neural message-passing baseline and show that this model, which incorporates a relational inductive bias, is superior at combinatorial generalization compared to a traditional recurrent neural network approach.

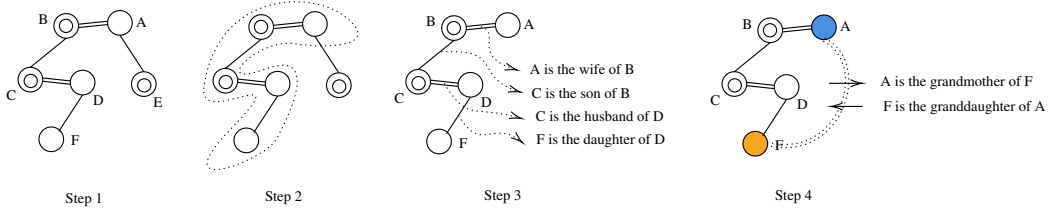
## 1 Introduction

Neural language understanding systems have been extremely successful at information extraction tasks, such as question answering (QA). An array of existing datasets are available, which test a system’s ability to extract factual answers text [1–5], as well as datasets that emphasize simple, commonsense inference (e.g., entailment between sentences) [6, 7]. However, it is difficult to evaluate a model’s reasoning ability in isolation using existing datasets. Most datasets combine several challenges of language processing into one, such as co-reference / entity resolution, incorporating world knowledge, and semantic parsing. Moreover, the state-of-the-art on all these existing benchmarks relies heavily on large, pre-trained language models [8, 9], highlighting that the primary difficulty in these datasets is incorporating the statistics of natural language, rather than reasoning.

In this work, we seek to directly evaluate and innovate on the compositional reasoning ability of a QA system. Inspired by CLEVR [10]—a synthetic computer vision dataset that isolates the challenges of relational reasoning—we propose a text based dataset for Compositional Language Understanding with Text-based Relational Reasoning (CLUTRR). Our initial version, CLUTTR v0.1, requires reasoning and generalizing about kinship relationships, and we plan to use our proposed data generation pipeline to extend the set of tasks in the future. We develop and evaluate strong baselines on CLUTTR v0.1, including a recurrent LSTM model and a message-passing graph neural network

---

\*Work done while being an intern at Samsung Advanced Institute of Technology (SAIT), Montreal



**Figure 1:** Data generation pipeline. Step 1: generate a kinship graph. Step 2: sample a path. Step 3: generate story by describing individual edges. Step 4: Predict relation between the first and the last nodes of the path.

(GNN). Our results highlight that the GNN, which incorporates a strong relational inductive bias, outperforms the LSTM at tasks requiring combinatorial generalization.

## 2 The CLUTRR dataset

To move away from fact-based extractive Q&A towards more relational reasoning, we consider the classic game of deducing family relations from text. Family relations have been used extensively in automated Knowledge Base (KB) completion tasks [11–14]. However, all of these systems operate on (sub)symbolic representations and partial KBs, i.e. they represent an entity *and* a relation with predicate symbols and are provided with a partial KBs. In this work, we instead learn directly from the textual descriptions, and we perform Cloze-style anonymization of entities. Thus, in our setting the model is not provided with a KB, and, instead, it must learn to *create* an (implicit) knowledge base as it deems fit. This setup is inspired both the CLEVR task [10] as well as the BAbI tasks [15]. Like BAbI, we emphasize reasoning in a controlled setting. However, going beyond the BAbI tasks, we introduce large amounts of distractor information, longer reasoning chains, and the primary focus of our dataset is that it explicitly requires combinatorial generalization, e.g., by training models on examples with  $k$  supporting facts and testing on examples that require  $> k$  supporting facts.

**Dataset construction.** The core idea behind the CLUTRR task is the following: given a text-based story describing a subset of a kinship graph, the goal is to predict the relationship between two entities, whose relationship is *not* stated in the story. Figure 1 illustrates the data generation process. Each example story,  $S$ , is a sequence of simple sentences, constructed in 3 steps:

1. Generate a random kinship graph  $G$  with  $m$  nodes. The nodes in the graph are the members of the *family* being represented and the edges represents relationships between the members. The relations belong to the relation set  $R = \{father, mother, son, daughter, husband, wife, grandfather, grandmother, grandson, granddaughter, brother, sister, father-in-law, mother-in-law, son-in-law\}$ .
2. Sample a simple path of  $k$  edges from this family graph.
3. For each edge in the sampled path, a template-based natural language description of the relation is sampled (e.g., “A is the wife of B”) and added to the story (with 2-3 possible templates per relation). Additionally, for each of the  $k + 1$  nodes in the path, 8 “distractor” sentences are sampled based on a set of distractor attributes (e.g., “A works at Y”) and randomly inserted in the story.

Given a generated story,  $S$ , the goal is to predict the relationship between the first entity  $e_0$  and the last entity  $e_k$  in the sampled path that was used to generate the story. Thus, the final task is  $|R|$ -way supervised classification, and to predict the correct relation, a model must infer the implicit knowledge graph represented by the story, reason about missing relationships in this implicit knowledge graph, and learn about logical regularities in kinship graphs (e.g., the child of a child is a grandchild). The Appendix contains further details on the graph generation pseudocode, natural language templates, distractor sentences, and sampling procedures.

**CLUTRR v0.1.** Our framework is highly extensible and modular, and in this work we focus on a prototype version, termed CLUTRR v0.1<sup>2</sup>, where we use the above methodology to sample 5000 stories for each possible path length  $k$ , ranging from  $k = 3$  to  $k = 6$ . We denote the sampled subsets for different paths lengths as  $M_1, \dots, M_6$ , and for each of these subsets we use 4000 examples for training and 1000 for testing. Separating the dataset into subsets based on the length of the underlying

<sup>2</sup><https://github.com/koustuvsinha/clutrr>

reasoning path allows us to explicitly evaluate how well models can generalize (e.g., training on  $M_k$  and testing on  $M_{k+1}$ ). Future versions will add realistic natural language variation via crowdsourcing.

### 3 Model

We introduce two strong baseline models for the CLUTRR dataset: the first model is based on a LSTM recurrent neural network and does not incorporate any relational inductive bias. The second model is a novel graph neural network (GNN) [16] architecture that is specifically designed to capture the relational structure of the CLUTRR reasoning task.

#### 3.1 Setup

The input to the model is a story  $S = (s_1, s_2, \dots, s_n)$ , where  $s_i = (w_1^i, w_2^i, \dots, w_m^i)$  is a sentence consisting of  $m$  words, each represented by a  $d$ -dimensional embedding. A subset of these words are the entities  $\{e_1, e_2, \dots, e_k\}$  (with Cloze-style anonymization), where each entity denotes a node in the latent kinship graph. Each sentence describes a relation between a pair of entities, or a “distractor” attribute of one of the entities. The task is therefore to use the input sentence to predict the relation between a pair of entities  $(e_i, e_j)$  as a  $|R|$ -way classification problem (see Section 2).

To approach the above problem, we describe our baseline models in terms of three components: a *reader* module to read the text, a *processor* module to perform reasoning over the text, and a *classifier* module which takes a query entity tuple  $(e_i, e_j)$  and predicts the relation  $r_{(e_i, e_j)} \in R$ .

#### 3.2 LSTM baseline

Firstly, we use a bidirectional LSTM (bi-LSTM) baseline [17]. We run the bi-LSTM *reader* over the document  $S$  to get the final document representation  $h_t$  and the intermediate representations for each word  $h_{w_i}$ , which also contains the query entities  $e_i$  and  $e_j$ . Next, we use a two layer MLP as the *classifier* module which is provided the concatenation of the intermediate hidden states of the query entities as well as the final document representation:  $p = \text{softmax}(\text{MLP}([\vec{h}_t, \vec{h}_{e_i}, \vec{h}_{e_j}]))$ , where  $\vec{h}_t = [\vec{h}_t, \overleftarrow{h}_t]$ ;  $\vec{h}_{e_k} = [\vec{h}_{e_k}, \overleftarrow{h}_{e_k}]$  for  $k \in (i, j)$ .

#### 3.3 GNN Baseline

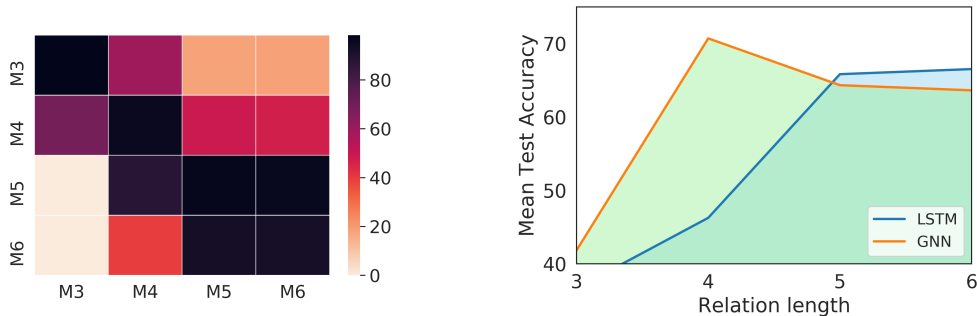
To incorporate a relational inductive bias, we use a graph neural network (GNN) model as a second strong baseline. In this model, the *reader* creates a unique edge representation from the relations by extracting the semantic information from the given text; the *processor* is a GNN model that computes node representations, and the *classifier* uses the node representations to classify the relation between the two query entities.

**Reader Module.** For simplicity, we assume that every sentence in a story describes a relationship between two family members, and we use the words occurring between entities to extract a learned edge embedding. The edge embedding for each node pair is calculated by pooling over the set of words occurring between them in the text:  $E_{i,j} = \text{pool}(\{w_1, w_2, \dots, w_p\})$ , where *pool* is a differentiable pooling function (e.g., max-pooling or an attention-weighted average). After computing the  $E_{i,j}$  embeddings from all sentences in the story, we obtain a graph  $G$  with edge features given by  $E_{i,j}$ .

**Processor Module.** The core component of the processor module is inherently a Message Passing Neural Network (MPNN) [16]. The input to this processor is a set of node features  $h = \{\vec{h}_1, \vec{h}_2, \dots, \vec{h}_K\}$ , where  $K$  is the maximum number of entities over the entire dataset. For each node, we calculate the incoming message from each of its neighbors in the subset  $\mathcal{N}(i)$  in the graph  $G$  as:

$$m_{i,j} = [E_{(i,j)}, \vec{p}_i, \vec{p}_j], \text{ where } j \in \mathcal{N}(i), \quad (1)$$

where  $\vec{p}_i$  is an embedding indicating the position of the node  $v_j$  within the graph  $G$ . The position information is represented as  $\vec{p}_i = [\vec{e}_i, \vec{g}]$  where  $\vec{e}_i$  is a random fixed embedding to represent the node, and  $\vec{g}$  is a parameter we learn to represent the graph structure. Intuitively, the parameter  $\vec{g}$  learns specific structural properties of the given graph and paired with a random embedding we can uniquely represent the node.



**Figure 2:** **Left:** the test accuracies for training data subset on the left axis and testing data subset on the right axis. **Right:** the mean testing accuracy as an area plot across the test data subsets.

**Table 1:** Generalizability experiment results

Training	$M_3$				$M_4$				$M_5$				$M_6$			
	$M_3$	$M_4$	$M_5$	$M_6$	$M_3$	$M_4$	$M_5$	$M_6$	$M_3$	$M_4$	$M_5$	$M_6$	$M_3$	$M_4$	$M_5$	$M_6$
LSTM	<b>100</b>	46.6	0	0	37	<b>100</b>	48	0	17.8	45.5	<b>100</b>	<b>100</b>	18.5	47.59	<b>100</b>	<b>100</b>
GNN	98.3	<b>69</b>	0	0	59.8	95.3	87.8	39.9	<b>18.4</b>	<b>49.4</b>	96.8	92.7	<b>18.5</b>	<b>47.6</b>	96.2	92.2
GNN-Attention	99.8	67.79	0.001	0	<b>65.7</b>	98.2	<b>89.9</b>	<b>44.1</b>	16.8	47.1	99.1	97.9	17.2	46.29	98.7	97.2

For each node, the incoming messages from its neighbors are averaged and combined together to form the final message:  $M_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} m_{i,j}$ . We also experimented with an attention style message combination similar to Graph Attention Networks (GAT)[18]. Specifically, the individual attention scores are calculated as  $a_{i,j} = W^\top \tanh(W[[\vec{h}_i, \vec{p}_i]; \vec{h}_j])$ , and the final attention weight is calculated as a softmax over the neighbors:  $\alpha_{i,j} = \frac{\exp(a_{i,j})}{\sum_{k \in \mathcal{N}(i)} \exp(a_{i,k})}$ . The combined message is a weighted sum over the attention:  $M_i = \sum_{j \in \mathcal{N}(i)} m_{i,j} \alpha_{i,j}$

Finally, the node representation is updated based on an LSTM update function similar to [19]:

$$\vec{h}_i = \text{LSTM}(M_i, (\vec{h}_i, c_i)) \quad (2)$$

**Classifier module.** The above message passing and update function is performed over  $r$  iterations, where  $r$  is a hyperparameter. After the  $r$  message-passing iterations, the final representations of each node are averaged to get the final graph representation  $H_T$ . The final relation classifier takes a concatenation of  $H_T$  and the node representations of the query entities:  $p = \text{softmax}(\text{MLP}([H_T, H_{v_i}, H_{v_j}]))$ .

## 4 Experimental Results

**Setup.** In order to explicitly evaluate the models’ performance on combinatorial generalization, we train on stories  $M_k$  generated from relation paths of length  $k$  and test on all the data subsets  $M_j$  for  $j \in k$ , i.e., we test how well the models can generalize from stories bases on varying lengths of relation paths (see Section 2 for further dataset details). We used 100-dimensional word embeddings for both the models, which are randomly initialized and fine-tuned. We treat the entities as Cloze style placeholders so we invalidate the embedding for the entity words after each training iteration. For the LSTM baseline we choose a 2 layer Bi-LSTM with 50 hidden dimensions (i.e., the full intermediate hidden states are 100-dimensional). For the GNN baseline we set the individual node embeddings to have 100 dimensions. We set the position embedding  $\vec{p}$  to have 15 dimensions where  $\vec{e}_i$  has 5 dimensions and  $\vec{g}$  has 10 dimensions. The message passing is run for  $r = 6$  iterations. We train both models using Adam optimizer with a learning rate of 0.001.

**Results.** From the results (Table 1) we can see that the GNN baseline model outperforms the LSTM significantly in terms of generalizaion across data subsets. However, we also note that GNN model does not get perfect scores when tested on the same data subset. This indicates that the LSTM is extremely powerful at rote learning patterns, but fails to generalize. In contrast, the GNN baseline picks up the compositional elements to build a relation from smaller relation chains ( $k \in (3, 4)$ ) and generalizes it into larger relation chains ( $k \in (5, 6)$ ) with significantly higher accuracy. This is

attributed to the inherent mechanism of relational reasoning with a GNN architecture. We further illustrate the generalization capability of the GNN model in Figure 2.

## 5 Conclusion

We present a new dataset and highly extensible data generation approach to evaluate relational reasoning on language, providing a focused evaluation of a model’s capacity for combinatorial generalization. We also present two strong baseline algorithms and show that a model with a strong relational inductive bias achieves superior generalization performance. As future work, we would like to extend our dataset and data generation approach by increasing the complexity of the relation set and adding crowdsourced (e.g., paraphrased) text for more natural language variation.

## Acknowledgement

The authors would like to acknowledge Sanghyun Yoo, Byung In Yoo, Jehun Jeon and Young-Seok Kim from Samsung Advanced Institute of Technology (SAIT), Korea for their helpful comments, feedback and discussion. The first author would also like to thank the entire On-Device Language Learning team from SAIT, Korea led by Dr Young Sang Choi for hosting him in the Summer of 2018 at Suwon, Korea.

## References

- [1] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 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, 2016.
- [2] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*, 2016.
- [3] Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. NewsQA: A machine comprehension dataset. *arXiv preprint*, pages 1–12, 2016. doi: 10.1016/B978-0-12-800077-9/00020-7.
- [4] Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. A corpus and cloze evaluation for deeper understanding of commonsense stories. In *Proceedings of NAACL-HLT*, pages 839–849, 2016.
- [5] Yu Su, Huan Sun, Brian Sadler, Mudhakar Srivatsa, Izzeddin Gur, Zenghui Yan, and Xifeng Yan. On generating characteristic-rich question sets for qa evaluation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 562–572, 2016.
- [6] Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*, 2015.
- [7] Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 1112–1122, 2018.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [9] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *NAACL*, 2018.

- [10] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 1988–1997. IEEE, 2017.
- [11] Stanley Kok and Pedro Domingos. Statistical predicate invention. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07*, pages 433–440, New York, NY, USA, 2007. ACM. ISBN 9781595937933. doi: 10.1145/1273496.1273551.
- [12] Stephen Muggleton. Inductive logic programming. *New Generation Computing*, 8(4):295–318, February 1991. ISSN 1882-7055. doi: 10.1007/BF03037089.
- [13] Nada Lavrac and Saso Dzeroski. Inductive logic programming. In *WLP*, pages 146–160. Springer, 1994.
- [14] Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, pages 3788–3800, 2017.
- [15] Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. Towards AI-Complete question answering: A set of prerequisite toy tasks. 2015. ISSN 0378-7753. doi: 10.1016/j.jpowsour.2014.09.131.
- [16] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272, 2017.
- [17] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- [18] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [19] Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. A graph-to-sequence model for amr-to-text generation. *arXiv preprint arXiv:1805.02473*, 2018.

## A Graph Generation

As described in Section 2, we perform the graph generation over 3 steps: generation of random kinship graph, sampling of a simple path of  $k$  edges, and replacing the edges with template-based natural language of the relation. The detailed pseudocode of data generation process is given in Algorithm 1. Specifically, we use two natural language template dictionaries:  $TD$  for kinship relation dictionary and  $AD$  for attribute relations. From  $TD$  for each node pair and their relation from the set  $R$  (refer Section 2), we sample a natural language description. For example, if  $n_a$  and  $n_b$  has a relation *significant other*, then one possible sampled template would be “A is the wife of B” depending on whether the gender of  $n_a$  is female.  $TD$  consists of several such templates in both gender relations for each  $r \in R$ .

### A.1 Distractor relations

For each node, we sample  $d$  distractor *attributes*. These attributes can range from where the person works, which school they are alumni of, to political preferences (Republican / Democrat). We choose a set of 8 such distractor relations  $D = \{works\_at, alumni\_of, school, location\_born, preferred\_social\_media, hobby, sport, political\_view\}$ . The dictionary  $AD$  contains templates for each of these distractor attributes which we sample as given in the Algorithm 1.

### A.2 Relation combinations

We sample relation path of length  $k$  from the graph(s). Depending on the number of paths extracted from a graph, we generate more graphs if necessary to match the number of training-testing rows. Each row in the data, as described in TEMPLATE STORY GENERATION procedure of Algorithm 1,

**Table 2:** Total unique relation paths for  $G = \{l_{max} = 3, c_{min} = 3, c_{max} = 3\}$ 

Path length	Unique paths	Unique paths with gender
3	20	40
4	84	168
5	305	610
6	978	1956
7	2814	5628

consists of a pair of relation path of length  $k$  and the target relation which is the relation among the first and the last node of the relation path. Each relation path does not contain duplicate nodes so as to remove cycles from the path. For example, a relation path of length 3 can have the following sequence:  $A \xrightarrow{child} B \xrightarrow{SO} C \xrightarrow{child} D$  (where  $SO$  denotes *significant other*). In this example, the target relation would be  $(A, D) \rightarrow grand[mother/father]$  depending on the gender of  $A$ . Note that in a target relation tuple we always refer the relation from the point of view of the first element. Thus if we inverse the above relation then the target changes :  $(D, A) \rightarrow grand[son/daughter]$ .

As we show in the example above, for each unique path we have two possible paths depending on the gender. This setup thus allows us to sample from the full possible combinations of kinship relations for each path length  $k$ . For example, total number of combinations of kinship relations for a graph of three levels and three children for each node is given in Table 2.

---

**Algorithm 1** CLUTRR story generation process

---

```
1: procedure GRAPH GENERATION
2:   Initialize a graph  $f$ 
3:   Add a node  $n$  to the graph  $f$  which serves as head of the family
4:   add the node  $n$  in parents array  $ps$ 
5:   for  $l = 0$  to  $l_{max}$  do  $\triangleright l_{max}$  is the maximum number of levels, parameter
6:     for  $p$  in  $[ps_0, \dots, ps_{last}]$  do:
7:       Add a partner node  $n_p$  of the opposite sex to the graph  $f$ 
8:       Add a significant other edge between  $p$  and  $n_p$ 
9:       Randomly determine the number of children  $c_r$  of  $(p, n_p)$  from a minimum  $c_{min}$  and
maximum  $c_{max}$ 
10:      Add  $|c|$  number of child nodes with random gender in the graph  $f$ 
11:      for  $c = 0$  to  $c_r$  do:
12:        Add a child edge between  $(c, p)$ 
13:        Add a child edge between  $(c, n_p)$ 
14:      Replace parents array  $ps$  with all children generated
15:   return graph  $f$ 
16: procedure PATH SAMPLING( $f, k$ )  $\triangleright k$  is the length of the relation we want to sample paths
17:   extract all siblings by grouping the nodes which have the same parent
18:   for each node pair  $(n_a, n_b)$  in the set of nodes  $N$  of the graph  $f$  do
19:     if there exists no edge between  $n_a$  and  $n_b$  then
20:       if there exists a shortest path between  $n_a$  and  $n_b$  then
21:         deduce the relation  $r_{(n_a, n_b)}$  among the shortest path from the set of rules  $R$ 
22:         Add an edge between  $(n_a, n_b)$  with this new relation
23:    $P \leftarrow$  empty array
24:   for each node pair  $(n_a, n_b)$  in the set of nodes  $N$  of the graph  $f$  do
25:     Extract all simple paths  $sp$  between the nodes  $n_a$  and  $n_b$ 
26:     Sample all paths of length  $k$ 
27:     add the paths to  $P$ 
28:   Return all paths  $P$ 
28: procedure TEMPLATE STORY GENERATION( $P$ )
29:   Initialize story-abstract pair  $SP$  array
30:   for each path  $p_i \in P$  do
31:     Initialize empty string for story  $S$ 
32:     Initialize empty string for abstract  $A$ 
33:     for each sequential node pair  $(n_a, n_b) \in p_i$  do
34:       Extract relation  $r_{(n_a, n_b)}$  from the edge between the nodes  $(n_a, n_b)$ 
35:       For the relation and the gender of  $n_b$ , sample from the template dictionary  $TD$  a
placeholder sentence  $t_p$ 
36:       Replace the entity placeholders with the name of  $r_{(n_a, n_b)}$  in  $t_p$ 
37:        $S \leftarrow S + t_p$ 
38:       Sample  $d$  attributes from each of the nodes  $(n_a, n_b)$ 
39:       for each attribute  $d_i \in d$  do
40:         Sample a template from attribute dictionary  $AD$  which is  $td$ .
41:         Replace the template with the node name and attribute name.
42:          $S \leftarrow S + td$ 
43:       Extract relation  $r_{(n_{P[0]}, n_{P[-1]})}$  from the edge between the nodes  $(n_{P[0]}, n_{P[-1]})$  which
are the first and the last nodes of the path  $P$ 
44:       Similarly, sample from the template dictionary  $TD$  a placeholder sentence  $t_p$  for the
relation and gender of  $n_{P[-1]}$ .
45:        $A = t_p$ 
46:        $SP \leftarrow SP + (S, A)$ 
Return  $SP$ 
```

---