

CycleGT: Unsupervised Graph-to-Text and Text-to-Graph Generation via Cycle Training

Qipeng Guo*[†]
Fudan University
qpguo16@fudan.edu.cn

Zhijing Jin*
Amazon Shanghai AI Lab
zjin@tue.mpg.de

Xipeng Qiu
Fudan University
xpqiu@fudan.edu.cn

Weinan Zhang
Shanghai Jiao Tong University
wnzhang@sjtu.edu.cn

David Wipf
Amazon Shanghai AI Lab
daviwipf@amazon.com

Zheng Zhang
Amazon Shanghai AI Lab
zhaz@amazon.com

Abstract

Two important tasks at the intersection of knowledge graphs and natural language processing are graph-to-text (G2T) and text-to-graph (T2G) conversion. Due to the difficulty and high cost of data collection, the supervised data available in the two fields are usually on the magnitude of tens of thousands, for example, 18K in the WebNLG 2017 dataset after preprocessing, which is far fewer than the millions of data for other tasks such as machine translation. Consequently, deep learning models for G2T and T2G suffer largely from scarce training data. We present CycleGT, an unsupervised training method that can bootstrap from fully non-parallel graph and text data, and iteratively back translate between the two forms. Experiments on WebNLG datasets show that our unsupervised model trained on the same number of data achieves performance on par with several fully supervised models. Further experiments on the non-parallel GenWiki dataset verify that our method performs the best among unsupervised baselines. This validates our framework as an effective approach to overcome the data scarcity problem in the fields of G2T and T2G.¹

1 Introduction

Knowledge graphs are a popular form of knowledge representation and central to many critical natural language processing (NLP) applications. One of the most important tasks, graph-to-text (G2T), aims to produce descriptive text that verbalizes the graphical data. For example, the knowledge graph triplet “(Allen Forest, genre, hip hop), (Allen Forest, birth year, 1981)” can be verbalized as “Allen Forest, a hip hop musician, was born in 1981.”

*Equal contribution.

[†]Work done during internship at Amazon Shanghai AI Lab.

¹Our code is available at <https://github.com/QipengGuo/CycleGT>.

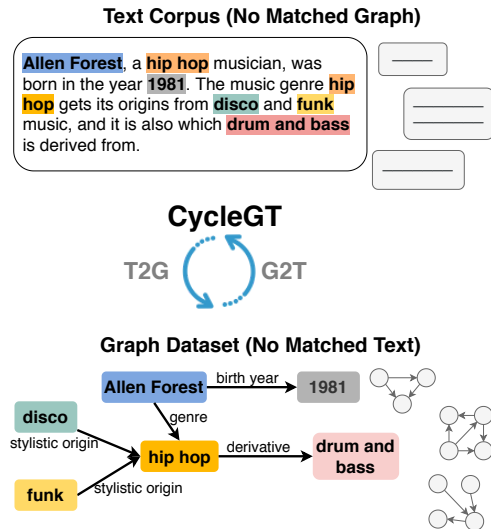


Figure 1: Given a text corpus, and a graph dataset, and no parallel (text, graph) pairs, our model CycleGT aims to jointly learn T2G and G2T in a cycle framework.

This has wide real-world applications, for instance, when a digital assistant needs to translate some structured information (e.g., the properties of the restaurant) to the human user. Another important task, text-to-graph (T2G), is to extract structures in the form of knowledge graphs from the text, so that all entities become nodes, and the relationships among entities form edges. This can serve many downstream tasks, such as information retrieval and reasoning. The two tasks can be seen as a dual problem, as shown in Figure 1.

However, most previous work has treated G2T and T2G as two separate supervised learning problems, for which the data annotation is very expensive. Therefore, both fields face the challenge of scarce parallel data. All current datasets are of a much smaller size than what is required to train the model to human-level performance. For example, the benchmark dataset WebNLG 2017 only has 18K text-graph pairs for training (after pre-

processing by (Moryossef et al., 2019)), which is several magnitudes fewer than the millions of data for neural machine translation (NMT) systems, such as the 4.4M paired sentences in the WMT14 dataset for NMT (Luong et al., 2015). As a result, most previous G2T and T2G models, which have to be trained on small datasets, display limited performance on both tasks (Koncel-Kedziorski et al., 2019; Moryossef et al., 2019; Luan et al., 2018).

To circumvent the limitations of scarce supervised data, we formulate both tasks in a cycle training framework, and also in the unsupervised manner with fully non-parallel datasets (as shown in Figure 1). The technical difficulty lies in the different modality of text and graphs, which can be intractable in a joint learning setting. We contribute an effective learning framework, CycleGT, which is iteratively trained with two cycle losses.

We first validate CycleGT on widely used WebNLG datasets, both WebNLG 2017 (Gardent et al., 2017b) and WebNLG+ 2020 (Castro-Ferreira et al., 2020) which is a more challenging extension of the 2017 dataset. Here CycleGT achieves performance that is comparable to many supervised baselines. Additionally, we also evaluate our model on the newly released GenWiki dataset (Jin et al., 2020b) with 1.3M non-parallel text and knowledge graphs. Our model is the best out of all unsupervised baselines, improving by +11.47 BLEU over the best baseline on GenWiki_{FINE} data and +6.26 BLEU on GenWiki_{FULL}.

The surprisingly high performance of CycleGT indicates that it is an effective approach to address the data scarcity problem in the fields of both G2T and T2G. Consequently, CycleGT can help pave the way for scalable, unsupervised learning, and benefit future research in both fields.

2 Formulation and Notations

For unsupervised graph-to-text and text-to-graph generation, we assume two non-parallel datasets:

- A text corpus $D_T = \{t_i\}_{i=1}^N$ consisting of N text sequences, and
- A graph dataset $D_G = \{g_j\}_{j=1}^M$ consisting of M graphs.

The constraint is that the graphs and text contain the same distribution of latent content z , but are different forms of surface realizations. Their marginal log-likelihood can be formulated with the shared

latent content z :

$$\log p(g) = \log \int_z p(g | z)p(z)dz, \quad (1)$$

$$\log p(t) = \log \int_z p(t | z)p(z)dz. \quad (2)$$

Our goal is to train two models in an unsupervised manner: G2T that generates text based on the graph, and T2G that produces the graph based on text. In this vein, there is an implicit assumption that a one-to-one mapping exists, at least approximately, between text and graphs.

Denote the parameters of G2T as θ , and parameters of T2G as φ . Now suppose there exists an *unseen* ground truth distribution $\mathcal{D}_{\text{Pair}}$ (e.g., as in the test set), where each paired text and graph (t, g) share the same content. Then the ideal objective is to maximize the log-likelihood of θ and φ over the text and graph pairs $(t, g) \sim \mathcal{D}_{\text{Pair}}$:

$$\begin{aligned} \mathcal{J}(\theta, \varphi) &= \mathbb{E}_{(t,g) \sim \mathcal{D}_{\text{Pair}}} [\log p(t | g; \theta) + \log p(g | t; \varphi)]. \end{aligned} \quad (3)$$

The major challenge of our task is that the ground truth distribution $\mathcal{D}_{\text{Pair}}$ is not available as training data. Instead, only the text corpus D_T and graph dataset D_G are observed separately without alignment. So our solution is to approximate the learning objective in Eq. (3) through non-parallel text and graph data. The resulting method we derive, namely CycleGT, will be introduced next. Note that we have also recently adopted the underlying CycleGT architecture in the context of developing an alternative conditional variational autoencoder model for handling non-bijective mappings (Guo et al., 2020b).

3 CycleGT Development

In this section, we will first introduce the G2T and T2G components in Section 3.1 and 3.2, respectively, and then discuss the iterative back translation training strategy of CycleGT in Section 3.3.

3.1 G2T Component

The model G2T : $\mathcal{G} \rightarrow \mathcal{T}$ takes as input a graph g and generates a text sequence \hat{t} that is a sufficient description of the graph. As pretrained models are shown to be very effective on G2T tasks (Ribeiro et al., 2020; Kale, 2020), we use T5 (Raffel et al., 2020), a large pretrained sequence-to-sequence model as the G2T component. Note that

the model architecture of G2T is flexible in our cycle training framework, and it can also be substituted by alternatives such as the GNN-LSTM architecture as proposed in [Koncel-Kedziorski et al. \(2019\)](#).

Given a knowledge graph g , we first linearize it to a sequence $\text{Seq}(g)$, using $\langle H \rangle$, $\langle R \rangle$, $\langle T \rangle$, to denote the head, relation, and tail of a triple in the graph. As our paper mainly addresses graphs that can be verbalized in several sentences, the concerned graphs are small-scale. Hence, we linearize the graph according to the given order of triplets in the WebNLG dataset, as different linearization orders of the graphs are not a major concern, and can be handled well by the T5 model.

The finetuning of T5 aims to find the optimal parameters θ^* to correctly encode the graph sequence and decode it to the sentence. For this purpose we use maximum likelihood estimation objective

$$\theta^* = \underset{\theta}{\operatorname{argmax}} \prod_{(t,g) \sim \mathcal{D}} p(t \mid \text{Seq}(g); \theta). \quad (4)$$

3.2 T2G Component

We then introduce the other component, the T2G model. The function $\text{T2G} : \mathcal{T} \rightarrow \mathcal{G}$ takes as input a text sequence t and extracts its corresponding graph \hat{g} , whose nodes are the entities and edges are the relations between two entities. As generating both the entities and relations of the graph in a differentiable way is generally intractable, we directly use the entities if they are given, or if not given, we use an off-the-shelf entity extraction model ([Qi et al., 2020](#)) that identifies all entities in the text with high accuracy. We denote the set of entities extracted from the text as $\text{NER}(t)$. We then predict the relations between every two entities to form the edges in the graph. Note that our T2G component presented below is also flexible and can be replaced with other models.

Proceeding further, we first obtain the embeddings of every entity $v_i \in \text{NER}(t)$ by average pooling the contextualized embedding of each word w_j in the entity term. This leads to

$$v_i = \frac{1}{\text{Len}(v_i)} \sum_{w_j \in v_i} \text{emb}(w_j), \quad (5)$$

$$\text{emb}(w_j) = \text{enc}(w_j, w_{<j}, w_{>j}),$$

where enc encodes the embedding of w_j by its preceding context $w_{<j}$ and succeeding content $w_{>j}$.

Based on the entity embeddings, we derive each edge of the graph using a multi-label classification

layer C . C takes in the two vertices of the edge and predicts the edge type, which includes the ‘‘relation’’ type, and the set of possible relations of entities, leading to

$$e_{ij} = C(v_i, v_j). \quad (6)$$

T2G training aims to find the optimal parameters φ^* that correctly encodes the text and predicts the graph via the maximum likelihood estimation objective

$$\begin{aligned} \varphi^* &= \underset{\varphi}{\operatorname{argmax}} \prod_{(t,g) \sim \mathcal{D}} p(g \mid t; \varphi) \\ &= \underset{\varphi}{\operatorname{argmax}} \prod_{(t,g) \sim \mathcal{D}} \prod_{i=0}^K \prod_{j=0}^K p(e_{ij} \mid v_i, v_j, t; \varphi), \end{aligned} \quad (7)$$

where $K = |\text{NER}(t)|$ is the number of entities in text t .

3.3 Cycle Training with Iterative Back Translation

In NLP, back translation ([Sennrich et al., 2016a](#)) is first proposed for machine translation where a sentence in the source language (e.g., English) should be able to first translate to the target language (e.g., French) and then again translated ‘‘back’’ to the original source language. The source sentence and the ‘‘back-translated’’ sentence should be aligned to be the same.

The essence of back translation is that a variable x and a bijective mapping function $f(\cdot)$ should satisfy $x = f^{-1}(f(x))$, where f^{-1} is the inverse function of f . In our case, G2T and T2G are inverse functions of each other, because one transforms the graph to text and the other converts text to the graph. Specifically, we align each text with its back-translated version, and also each graph with its back-translated version via the objectives

$$\mathcal{L}_{\text{CycT}}(\theta) = \mathbb{E}_{t \in D_T} [-\log p(t \mid \text{T2G}_\varphi(t); \theta)], \quad (8)$$

$$\mathcal{L}_{\text{CycG}}(\varphi) = \mathbb{E}_{g \in D_G} [-\log p(g \mid \text{G2T}_\theta(g); \varphi)]. \quad (9)$$

An equivalent way to interpret Eqs. (8) and (9) is that due to the unavailability of paired text and graphs, we approximate $\mathcal{D}_{\text{Pair}}$ with \hat{D}_{Pair} , a synthetic set of text-graph pairs generated by the two models. \hat{D}_{Pair} consists of $(t, \text{T2G}_\varphi(t))$ and $(\text{G2T}_\theta(g), g)$ for every $t \in D_T, g \in D_G$, leading to

$$\begin{aligned}\mathcal{L}_{\text{CycT}}(\theta) &= \mathbb{E}_{t \in D_T}[-\log p(t \mid \text{T2G}_\varphi(t); \theta)] \\ &= \mathbb{E}_{(t, \hat{g}) \in \hat{D}_{\text{Pair}}}[-\log p(t \mid \hat{g}; \theta)],\end{aligned}\tag{10}$$

$$\begin{aligned}\mathcal{L}_{\text{CycG}}(\varphi) &= \mathbb{E}_{g \in D_G}[-\log p(g \mid \text{G2T}_\theta(t); \varphi)] \\ &= \mathbb{E}_{(\hat{t}, g) \in \hat{D}_{\text{Pair}}}[-\log p(g \mid \hat{t}; \varphi)].\end{aligned}\tag{11}$$

As such, to the extent that these approximations are accurate, the sum of Eqs. (8) and (9) reasonably approximates the log likelihoods from Eq. (3).

Note that $\mathcal{J}(\theta, \varphi) = \mathcal{L}_{\text{CycT}}(\theta) + \mathcal{L}_{\text{CycG}}(\varphi)$ holds when \hat{D}_{Pair} has the same distribution as D_{Pair} . In our framework, we iteratively improve the G2T and T2G models using an iterative back translation (IBT) training scheme (Hoang et al., 2018), with the goal of reducing the discrepancy between the distribution of D_{Pair} and \hat{D}_{Pair} . Specifically, we repeatedly alternate the optimization of the two cycles described by Eqs. (10) and (11) over the corresponding θ or φ .

3.4 Challenges of Cycle Training for G2T and T2G

One of the challenges of our problem-specific use of IBT, even with the two established cycle losses above, is the *non-differentiability*, which constitutes a fundamental difference between our model and the line of work represented by CycleGAN (Zhu et al., 2017). For our cycle losses $\mathcal{L}_{\text{CycT}}$ and $\mathcal{L}_{\text{CycG}}$, the intermediate model outputs are *non-differentiable*. For example, in the G-Cycle (graph→text→graph), the intermediate text is decoded in a discrete form to natural language. Hence, the graph-to-text part G2T_θ will not be differentiable, and the final loss can only be propagated to the latter part, text-to-graph T2G_φ . Hence, when alternatively optimizing the two cycle losses, we first fix φ to optimize θ for the text cycle $\mathcal{L}_{\text{CycT}}$, and then fix θ to optimize φ for the graph cycle $\mathcal{L}_{\text{CycG}}$.

Although an analogous non-differentiability issue is shared by unsupervised NMT works (Lample et al., 2018; Artetxe et al., 2018), these approaches rely on other regularization factors such as autoencoder losses, adversarial losses, and warm start strategies. In contrast, our streamlined approach CycleGT relies solely on cycle training and the inductive biases of the T2G and G2T modules.

4 Experiments

4.1 Datasets

WebNLG 2017 Dataset Our first experiment uses the WebNLG 2017 dataset (Gardent et al., 2017b), which is widely used for graph-to-text generation.² Each graph consists of 2 to 7 triplets extracted from DBpedia, and the text is collected by asking crowd-source workers to describe the graphs. The dataset includes 10 categories in the training data: Airport, Astronaut, Building, City, ComicsCharacter, Food, Monument, SportsTeam, University, and WrittenWork. There are also 5 additional unseen categories in the test data: Athlete, Artist, CelestialBody, MeansOfTransportation, and Politician. We follow the preprocessing steps in Moryossef et al. (2019) to obtain the text-graph pairs (with entity annotation) for 13,036 training, 1,642 validation, and 4,928 test samples.

WebNLG 2020 Dataset We also use CycleGT to participate in the WebNLG 2020 challenge, which is an extension of the WebNLG 2017 dataset covering broader categories. Specifically, WebNLG 2020 includes all the 15 categories of WebNLG 2017 in the training set, along with a new category, Company. Its test data also add three additional unseen categories, Film, Scientist, and MusicalWork. Overall, the dataset has 35,426 training, 4,464 validation, and 1,779 test samples.

GenWiki Dataset Apart from the WebNLG datasets that have to be processed to fit non-parallel setting of our framework, we also evaluate with a natural non-parallel dataset, GenWiki (Jin et al., 2020b). GenWiki is formed from DBpedia knowledge graphs, and unpaired natural text in Wikipedia articles with the same topics as the knowledge graphs. The text has a larger vocabulary and more variety than the crowdsourced text in WebNLG. We report results on two variations of GenWiki, the full dataset GenWiki_{FULL}, and a fine, distantly aligned version GenWiki_{FINE}.

4.2 Evaluation Metrics

For G2T on WebNLG 2017 and GenWiki, we use the metrics that are mostly reported by other previous work (Moryossef et al., 2019) so that

²WebNLG 2017 is the most appropriate dataset, because in other datasets such as relation extraction datasets (Walker et al., 2006), the graph only contain a very small subset of the information in the text. It can be downloaded from https://webnlg-challenge.loria.fr/challenge_2017/.

we can have head-to-head comparison with existing systems. Specifically, we adopt the metrics from (Moryossef et al., 2019), i.e., BLEU (Papineni et al., 2002), Meteor (Banerjee and Lavie, 2005), ROUGE_L (Lin, 2004) and CIDEr (Vedantam et al., 2015), to measure the closeness of the reconstructed text (model output) to the input text.³ Briefly, they measure the n-gram precision, recall, or F-scores between the model outputs and the (ground-truth) references. In contrast, for T2G evaluation on WebNLG 2017, we use the micro and macro F1 scores of the relation types of the edges, following the standard practice in relation extraction (Miwa and Bansal, 2016a; Zhou et al., 2016).

Turning to WebNLG 2020, we report most major metrics listed on the leaderboard (Moussalem et al., 2020). For G2T, the metrics include BLEU, METEOR, chrF++ (Popovic, 2017), TER (Snoover et al., 2006), BERT_{F1} (Zhang et al., 2020), and BLUERT (Sellam et al., 2020).

4.3 Comparison Systems for WebNLG 2017

Unsupervised Baselines As cycle training models are unsupervised learning methods, we include the following unsupervised baselines:

- *RuleBased* (Schmitt et al., 2020): This baseline involves simply iterating through the graph and concatenating the text of each triplet.
- *Graph-Text Back Translator (GT-BT)* (Schmitt et al., 2020): This approach first serializes the graph and then applies a back translation model. Since the code for GT-BT is not yet available, we adopt their reported results as a reference.

Supervised Baselines We also compare with *supervised* systems using the original supervised training data. Since there is no existing work that jointly learns graph-to-text and text-to-graph in a supervised way, we can only use models that address one of the two tasks. For graph-to-text generation, we compare with

- *Melbourne*: The best supervised system submitted to the WebNLG challenge 2017 (Gardent et al., 2017b), which uses an encoder-decoder architecture with attention.
- *StrongNeural* (Moryossef et al., 2019): An

enhanced version of the common encoder-decoder model.

- *BestPlan* (Moryossef et al., 2019): A special entity ordering algorithm applied before neural text generation.
- *GraphWriter* (Koncel-Kedziorski et al., 2019): A graph attention network with LSTMs.
- *Seg&Align* (Shen et al., 2020): An approach that first segments the text into small units and then learns the alignment between data and target text segments.
- *Planner* (Zhao et al., 2020): A planner with relational graph convolutional networks and LSTMs.
- *T5-Large* (Kale, 2020): An application of the pretrained T5-Large model for serialized graph-to-text generation.
- *T5-Large* (Ribeiro et al., 2020): Another instance of the pretrained T5-Large model for serialized graph-to-text generation.
- *Supervised G2T* (our implementation) : Our implementation of T5-base, which is the same as the G2T component in our CycleGT.

For text-to-graph generation, we compare with state-of-the-art models including

- *OnePass* (Wang et al., 2019): A BERT-based relation extraction model.
- *T2G* (our implementation): The BiLSTM model that we adopt as the text-to-graph component in the cycle training of CycleGT.

4.4 Comparison Systems for GenWiki

Since GenWiki is a non-parallel dataset, we can only train unsupervised baselines or transfer models developed using other data. We compare against:

- *RuleBased* (Schmitt et al., 2020): Same as was applied for WebNLG 2017 evaluation (see above).
- *DirectTransfer*: A model trained on the supervised WebNLG 2017 dataset, and then deployed on the GenWiki test set.
- *NoisySupervised*: A baseline formed by first constructing distantly aligned pairs on the whole training data by entity overlap, and then learning from these distant supervised pairs with the G2T model.

Note that for fair comparison, we use GraphWriter (Koncel-Kedziorski et al., 2019) as the architecture for all G2T models including our CycleGT on GenWiki.

³We calculate all metrics using the pycocoevalcap tool (<https://github.com/salaniz/pycocoevalcap>).

	G2T Performance				T2G Performance	
	BLEU	METEOR	ROUGE _L	CIDEr	Micro F1	Macro F1
Supervised Models (G2T Only)						
Melbourne (introduced in (Gardent et al., 2017a))	45.0	0.376	63.5	2.81	–	–
StrongNeural (Moryossef et al., 2019)	46.5	0.392	65.4	2.87	–	–
BestPlan (Moryossef et al., 2019)	47.4	0.391	63.1	2.69	–	–
GraphWriter (Koncel-Kedziorski et al., 2019)	45.8	0.356	68.6	3.14	–	–
Seg&Align (Shen et al., 2020)	46.1	0.398	65.4	2.64	–	–
Planner (Zhao et al., 2020)	52.8	0.450	–	–	–	–
T5-Large (Kale, 2020)	57.1	0.440	–	–	–	–
T5-Large (Ribeiro et al., 2020)	59.3	0.440	–	–	–	–
Supervised G2T by T5-Base (Our Implementation)	56.4	0.445	69.1	3.86	–	–
Supervised Models (T2G Only)						
Supervised T2G (Our Implementation)	–	–	–	–	60.6	50.7
OnePass (Wang et al., 2019)	–	–	–	–	66.2	52.2
Unsupervised Models						
RuleBased (Schmitt et al., 2020)	18.3	0.336	–	–	0.0	0.0
GT-BT (Schmitt et al., 2020)	37.7	0.355	–	–	39.1	–
CycleGT (Unsup.)	55.5	0.437	68.3	3.81	58.4	46.4

Table 1: T2G and G2T performance of supervised and unsupervised models on WebNLG 2017 data.

	GenWiki _{FINE}				GenWiki _{FULL}			
	BLEU	METEOR	ROUGE _L	CIDEr	BLEU	METEOR	ROUGE _L	CIDEr
Rule-Based (Schmitt et al., 2020)	13.45	30.72	40.93	1.26	13.45	30.72	40.93	1.26
DirectTransfer (Jin et al., 2020b)	13.89	25.76	39.75	1.26	13.89	25.76	39.75	1.26
NoisySupervised (Jin et al., 2020b)	30.12	28.12	56.96	2.52	35.03	33.45	58.14	2.63
CycleGT	41.59	35.72	63.31	3.57	41.29	35.39	63.73	3.53

Table 2: The performance of our CycleGT and three unsupervised baselines on GenWiki_{FINE} and GenWiki_{FULL}.

4.5 Implementation Details

Processing WebNLG Data To test unsupervised models and baselines, we construct a non-parallel version of the training and validation sets by separating all the text in the dataset to form a text corpus, and all the graphs to build a graph dataset. We ensure that the order within the text and graph datasets are shuffled so that the data is fully non-parallel.

Model Details Our training framework can adapt to different T2G and G2T modules. For this study, we instantiate our CycleGT framework with the pre-trained T5 (Raffel et al., 2020) as the G2T model, and a BiLSTM for the T2G architecture, with two layers comprised of 512 hidden units each. We train the model for 30 epochs.

4.6 Main Results

WebNLG 2017 Results As shown in Table 1, for G2T generation, even with no pairing information between text and graphs, our unsupervised CycleGT model can achieve a 55.5 BLEU score, which is on par with the 56.4 BLEU score obtained using the same T2G model with T5-Base that we trained on supervised data. CycleGT also outperforms the unsupervised baseline RuleBased and

GT-BT (Schmitt et al., 2020) by a clear margin. Additionally, for the T2G task, CycleGT also achieves scores on par with the same T2G component but trained on the supervised data.

WebNLG 2020 Results We also submitted our system to the WebNLG 2020 Challenge. In Table 3, we list competing models from the competition leaderboard, as well as two baseline results published by the organizers. Since the challenge provides supervised training data, most (and possibly all) other competing systems are supervised models, such as the top-1 system, ID18 (Guo et al., 2020a), which finetunes T5 on the supervised graph-text pairs. As can be seen, our G2T performance is competitive with many methods even without supervision.

GenWiki Results We show the performance of CycleGT on the G2T task using GenWiki data in Table 2. Among the four models under consideration, the RuleBased and DirectTransfer methods do not perform well on the purely unsupervised dataset. The performance of DirectTransfer indicates that even though WebNLG and GenWiki are similar wiki-based datasets, it is difficult to make a model trained on one corpus to behave similarly

	BLEU	METEOR	CHRF++	TER	BLUERT
ID18	53.98	0.417	0.690	0.406	0.62
ID30	53.54	0.414	0.688	0.416	0.61
ID30_1	52.07	0.413	0.685	0.444	0.58
ID34	52.67	0.413	0.686	0.423	0.6
ID5	51.74	0.411	0.679	0.435	0.6
ID35	51.59	0.409	0.681	0.431	0.59
ID23	51.74	0.403	0.669	0.417	0.61
ID2	50.34	0.398	0.666	0.435	0.57
ID15	40.73	0.393	0.646	0.511	0.45
CycleGT (Unsup.)	44.56	0.387	0.637	0.479	0.54
ID12	40.29	0.386	0.634	0.504	0.45
ID11	39.84	0.384	0.632	0.517	0.43
ID4	50.93	0.384	0.636	0.454	0.54
ID26	50.43	0.382	0.637	0.439	0.57
Official Baseline1	40.57	0.373	0.621	0.517	0.47
ID17	39.55	0.372	0.613	0.536	0.37
ID31_2	41.03	0.367	0.608	0.522	0.39
Official Baseline2	37.89	0.364	0.606	0.553	0.42
ID21	31.98	0.350	0.545	0.629	0.4
ID13_11	38.37	0.343	0.584	0.587	0.33
ID14	39.12	0.337	0.579	0.564	0.37
ID13	38.2	0.335	0.571	0.577	0.29
ID13_3	39.19	0.334	0.569	0.565	0.28
ID13_4	38.85	0.332	0.569	0.573	0.3
ID13_2	38.01	0.331	0.565	0.583	0.27
ID13_6	37.96	0.331	0.566	0.580	0.28
ID10	22.84	0.326	0.534	0.696	-0.03
ID13_7	36.07	0.324	0.555	0.599	0.24
ID13_8	36.6	0.322	0.554	0.594	0.25
ID20	31.26	0.316	0.542	0.659	0.31
ID26_1	27.5	0.305	0.519	0.846	0.03
ID13_9	28.71	0.243	0.448	0.689	-0.09
ID13_10	30.79	0.239	0.439	0.677	-0.09
ID13_5	23.08	0.225	0.421	0.743	-0.19
ID17_1	24.45	0.223	0.425	0.739	-0.22

Table 3: Leaderboard of the submitted systems and two official baselines (Baseline1 and Baseline2) in the WebNLG 2020 English RDF-to-text challenge as ranked by METEOR. With all metrics, larger is better, with the exception of TER where lower is better.

on another slightly different one. The NoisySupervised model performs relatively well, scoring over 30 BLEU points. And our model CycleGT is the strongest, outperforming NoisySupervised by around 10 BLEU points.

5 Related Work

We will first give an overview of the fields of T2G and T2G separately, and then introduce the cycle learning.

Data-to-Text Generation As a classic problem in text generation (Kukich, 1983; McKeown, 1992), data-to-text generation aims to automatically produce text from structured data (Reiter and Dale, 1997; Liang et al., 2009). Due to the expensive collection, all the data-to-text datasets are very small, such as the 5K air travel dataset (Ratnaparkhi, 2000), 22K WeatherGov (Liang et al., 2009), 7K Robocup (Chen and Mooney, 2008), and 5K RotoWire on basketball games (Wiseman et al., 2017). As for the methodology, tra-

ditional approaches adopt a pipeline system (Kukich, 1983; McKeown, 1992) of content planning, sentence planning, and surface realization. Recent advances in neural networks give birth to end-to-end systems (Lebret et al., 2016; Wiseman et al., 2017; Koncel-Kedziorski et al., 2019; Kale, 2020) that does not use explicit planning but directly an encoder-decoder architecture (Bahdanau et al., 2015; Vaswani et al., 2017).

Relation Extraction Relation Extraction (RE) is the core problem in text-to-graph conversion, as its former step, entity recognition, have off-the-shelf tools with good performance (Lample et al., 2016; Qian et al., 2019; Straková et al., 2019; Clark et al., 2018; Jin et al., 2020c). RE aims to classify the relation of entities given a shared textual context. Conventional approaches hand-crafted lexical and syntactic features (Hendrickx et al., 2010; Rink and Harabagiu, 2010). With the recent advancement of deep neural networks, many models based on CNN (Zeng et al., 2014; dos Santos et al., 2015; Nguyen and Grishman, 2015), RNN (Socher et al., 2012; Zhang and Wang, 2015; Miwa and Bansal, 2016b; Zhou et al., 2016), and BERT (Wang et al., 2019) achieve high performance in many datasets. However, constrained by the small datasets of only several hundred or several thousand data points (Walker et al., 2006; Hendrickx et al., 2010; Gábor et al., 2018), recent research shifts to distant supervision than model innovation (Mintz et al., 2009; Zeng et al., 2015; Lin et al., 2016).

Cycle Training The concept of leveraging the transitivity of two functions inverse to each other has been widely observed on a variety of tasks. In computer vision, the forward-backward consistency has been used since last decade (Kalal et al., 2010; Sundaram et al., 2010), and training on cycle consistency has recently been extensively applied on image style transfer (Zhu et al., 2017; Godard et al., 2017). In language, back translation (Sennrich et al., 2016b; Edunov et al., 2018; Jin et al., 2020a) and dual learning (Cheng et al., 2016; He et al., 2016) have also been an active area of research centered on UMT. Similar techniques can also be seen on tasks such as language style transfer (Shen et al., 2017; Jin et al., 2019). Finally, cycle training has recently been applied in the context of text-graph conversion (Schmitt et al., 2020); we compare against this approach in Table 1 (see

GT-BT results).

6 Conclusion

We have developed a cycle learning framework for both text-to-graph and graph-to-text generation in an unsupervised way. Experimental results validate that the proposed model achieves comparable results to many supervised models, given the same amount of unsupervised data. Moreover, when comparing with other unsupervised models, our CycleGT model displays a clear advantage.

Acknowledgements

We thank colleagues at the Amazon Shanghai AI lab, including Xiangkun Hu, Hang Yan, and many others for insightful discussions that constructively helped this work.

References

- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. [Unsupervised neural machine translation](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: an automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL 2005, Ann Arbor, Michigan, USA, June 29, 2005*, pages 65–72. Association for Computational Linguistics.
- Thiago Castro-Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussalem, and Anastasia Shimorina. 2020. The 2020 bilingual, bi-directional webnlg+ shared task: Overview and evaluation results (webnlg+ 2020). In *Proceedings of the 3rd WebNLG Workshop on Natural Language Generation from the Semantic Web (WebNLG+ 2020)*, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- David L. Chen and Raymond J. Mooney. 2008. [Learning to sportscast: a test of grounded language acquisition](#). In *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008*, volume 307 of *ACM International Conference Proceeding Series*, pages 128–135. ACM.
- Yong Cheng, Wei Xu, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. [Semi-supervised learning for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc V. Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1914–1925. Association for Computational Linguistics.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 489–500. Association for Computational Linguistics.
- Kata Gábor, Davide Buscaldi, Anne-Kathrin Schumann, Behrang QasemiZadeh, Haifa Zargayouna, and Thierry Charnois. 2018. [Semeval-2018 task 7: Semantic relation extraction and classification in scientific papers](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2018, New Orleans, Louisiana, USA, June 5-6, 2018*, pages 679–688.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017a. [Creating training corpora for NLG micro-planners](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 179–188. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017b. [The webnlg challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation, INLG 2017, Santiago de Compostela, Spain, September 4-7, 2017*, pages 124–133. Association for Computational Linguistics.
- Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. 2017. [Unsupervised monocular depth estimation with left-right consistency](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6602–6611. IEEE Computer Society.
- Qipeng Guo, Zhijing Jin, Ning Dai, Xipeng Qiu, Xiangyang Xue, David Wipf, and Zheng Zhang. 2020a. P2: A plan-and-pretrain approach for knowledge graph-to-text generation. In *Proceedings of the 3rd WebNLG Workshop on Natural Language Generation from the Semantic Web (WebNLG+ 2020)*,

- Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Qipeng Guo, Zhijing Jin, Ziyu Wang, Xipeng Qiu, Weinan Zhang, Jun Zhu, Zheng Zhang, and David Wipf. 2020b. Fork or fail: Cycle-consistent training with many-to-one mappings. *under review*.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. [Dual learning for machine translation](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 820–828.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. [Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation, SemEval@ACL 2010, Uppsala University, Uppsala, Sweden, July 15-16, 2010*, pages 33–38.
- Cong Duy Vu Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. [Iterative back-translation for neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation, NMT@ACL 2018, Melbourne, Australia, July 20, 2018*, pages 18–24. Association for Computational Linguistics.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020a. [A simple baseline to semi-supervised domain adaptation for machine translation](#). *CoRR*, abs/2001.08140.
- Zhijing Jin, Qipeng Guo, Xipeng Qiu, and Zheng Zhang. 2020b. [Genwiki: A dataset of 1.3 million content-sharing text and graphs for unsupervised graph-to-text generation](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020*. Association for Computational Linguistics.
- Zhijing Jin, Di Jin, Jonas Mueller, Nicholas Matthews, and Enrico Santus. 2019. [Imat: Unsupervised text attribute transfer via iterative matching and translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3095–3107. Association for Computational Linguistics.
- Zhijing Jin, Yongyi Yang, Xipeng Qiu, and Zheng Zhang. 2020c. [Relation of the relations: A new paradigm of the relation extraction problem](#). *CoRR*, abs/2006.03719.
- Zdenek Kalal, Krystian Mikolajczyk, and Jiri Matas. 2010. [Forward-backward error: Automatic detection of tracking failures](#). In *20th International Conference on Pattern Recognition, ICPR 2010, Istanbul, Turkey, 23-26 August 2010*, pages 2756–2759. IEEE Computer Society.
- Mihir Kale. 2020. Text-to-text pre-training for data-to-text tasks. *arXiv preprint arXiv:2005.10433*.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. [Text generation from knowledge graphs with graph transformers](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2284–2293. Association for Computational Linguistics.
- Karen Kukich. 1983. [Design of a knowledge-based report generator](#). In *21st Annual Meeting of the Association for Computational Linguistics, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, June 15-17, 1983*, pages 145–150. ACL.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 260–270. The Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. [Unsupervised machine translation using monolingual corpora only](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Rémi Lebret, David Grangier, and Michael Auli. 2016. [Neural text generation from structured data with application to the biography domain](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1203–1213. The Association for Computational Linguistics.
- Percy Liang, Michael I. Jordan, and Dan Klein. 2009. [Learning semantic correspondences with less supervision](#). In *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore*, pages 91–99. The Association for Computer Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. [Neural relation extraction with selective attention over instances](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Yi Luan, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [The UWNLP system at semeval-2018 task 7: Neural relation extraction model with selectively incorporated concept embeddings](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2018, New Orleans, Louisiana, USA, June 5-6, 2018*, pages 788–792.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Kathleen R. McKeown. 1992. *Text generation - using discourse strategies and focus constraints to generate natural language text*. Studies in natural language processing. Cambridge University Press.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *ACL 2009, Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore*, pages 1003–1011. The Association for Computer Linguistics.
- Makoto Miwa and Mohit Bansal. 2016a. [End-to-end relation extraction using LSTMs on sequences and tree structures](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1105–1116, Berlin, Germany. Association for Computational Linguistics.
- Makoto Miwa and Mohit Bansal. 2016b. [End-to-end relation extraction using lstms on sequences and tree structures](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. [Step-by-step: Separating planning from realization in neural data-to-text generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2267–2277. Association for Computational Linguistics.
- Diego Moussalem, Paramjot Kaur, Thiago Castro-Ferreira, Chris van der Lee, Conrads Felix Shimorina, Anastasia, Michael Röder, René Speck, Claire Gardent, Simon Mille, Nikolai Ilinykh, and Axel-Cyrille Ngonga Ngomo. 2020. [A general benchmarking framework for text generation](#). In *Proceedings of the 3rd WebNLG Workshop on Natural Language Generation from the Semantic Web (WebNLG+ 2020)*, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2015. [Relation extraction: Perspective from convolutional neural networks](#). In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing, VS@NAACL-HLT 2015, June 5, 2015, Denver, Colorado, USA*, pages 39–48.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Maja Popovic. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation, WMT 2017, Copenhagen, Denmark, September 7-8, 2017*, pages 612–618. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). *CoRR*, abs/2003.07082.
- Yujie Qian, Enrico Santus, Zhijing Jin, Jiang Guo, and Regina Barzilay. 2019. [Graphie: A graph-based framework for information extraction](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 751–761. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Adwait Ratnaparkhi. 2000. [Trainable methods for surface natural language generation](#). In *6th Applied Natural Language Processing Conference, ANLP 2000, Seattle, Washington, USA, April 29 - May 4, 2000*, pages 194–201. ACL.
- Ehud Reiter and Robert Dale. 1997. [Building applied natural language generation systems](#). *Natural Language Engineering*, 3(1):57–87.
- Leonardo FR Ribeiro, Martin Schmitt, Hinrich Schütze, and Iryna Gurevych. 2020. [Investigating pretrained language models for graph-to-text generation](#). *arXiv preprint arXiv:2007.08426*.

- Bryan Rink and Sanda M. Harabagiu. 2010. [UTD: classifying semantic relations by combining lexical and semantic resources](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation, SemEval@ACL 2010, Uppsala University, Uppsala, Sweden, July 15-16, 2010*, pages 256–259.
- Cícero Nogueira dos Santos, Bing Xiang, and Bowen Zhou. 2015. [Classifying relations by ranking with convolutional neural networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 626–634.
- Martin Schmitt, Sahand Sharifzadeh, Volker Tresp, and Hinrich Schütze. 2020. An unsupervised joint system for text generation from knowledge graphs and semantic parsing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, Online. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur P. Parikh. 2020. [BLEURT: learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7881–7892. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi S. Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6830–6841.
- Xiaoyu Shen, Ernie Chang, Hui Su, Cheng Niu, and Dietrich Klakow. 2020. [Neural data-to-text generation via jointly learning the segmentation and correspondence](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, pages 7155–7165. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Lina Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of association for machine translation in the Americas*, volume 200. Cambridge, MA.
- Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. [Semantic compositionality through recursive matrix-vector spaces](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL 2012, July 12-14, 2012, Jeju Island, Korea*, pages 1201–1211.
- Jana Straková, Milan Straka, and Jan Hajic. 2019. [Neural architectures for nested NER through linearization](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5326–5331. Association for Computational Linguistics.
- Narayanan Sundaram, Thomas Brox, and Kurt Keutzer. 2010. [Dense point trajectories by gpu-accelerated large displacement optical flow](#). In *Computer Vision - ECCV 2010, 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part I*, volume 6311 of *Lecture Notes in Computer Science*, pages 438–451. Springer.
- 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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. [Cider: Consensus-based image description evaluation](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4566–4575. IEEE Computer Society.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57.
- Haoyu Wang, Ming Tan, Mo Yu, Shiyu Chang, Dakuo Wang, Kun Xu, Xiaoxiao Guo, and Saloni Potdar. 2019. [Extracting multiple-relations in one-pass with pre-trained transformers](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1371–1377.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. [Challenges in data-to-document generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2253–2263. Association for Computational Linguistics.

- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. [Distant supervision for relation extraction via piecewise convolutional neural networks](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1753–1762.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. [Relation classification via convolutional deep neural network](#). In *COLING 2014, 25th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, August 23-29, 2014, Dublin, Ireland*, pages 2335–2344.
- Dongxu Zhang and Dong Wang. 2015. [Relation classification via recurrent neural network](#). *CoRR*, abs/1508.01006.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Chao Zhao, Marilyn Walker, and Snigdha Chaturvedi. 2020. [Bridging the structural gap between encoding and decoding for data-to-text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2481–2491, Online. Association for Computational Linguistics.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. [Attention-based bidirectional long short-term memory networks for relation classification](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. [Unpaired image-to-image translation using cycle-consistent adversarial networks](#). In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2242–2251. IEEE Computer Society.