Artificial Intelligence and Human-Computer Interaction
Y. Ye and P. Siarry (Eds.)

doi:10.3233/FAIA240132

Dynamic Knowledge Graph-Based Dialogue Generation with Improved Adversarial Meta-Learning

Hongcai Xu and Junpeng Bao[1]
Xi'an Jiaotong University, China
ORCiD ID: Hongcai Xu https://orcid.org/0000-0001-7225-4446

Abstract. Knowledge graph-based dialogue systems are capable of generating more informative responses and can implement sophisticated reasoning mechanisms. However, these models do not take into account the sparseness and incompleteness of knowledge graph (KG) and cannot be applied to dynamic KG. This paper proposes a dynamic Knowledge graph-based dialogue generation method with improved adversarial Meta-Learning (ADML). ADML formulates dynamic knowledge triples as a problem of adversarial attack and incorporates the objective of quickly adapting to dynamic knowledge-aware dialogue generation. The model can initialize the parameters and adapt to previous unseen knowledge so that training can be quickly completed based on only a few knowledge triples. We show that our model significantly outperforms other baselines. We evaluate and demonstrate that our method adapts extremely fast and well to dynamic knowledge graph-based dialogue generation.

Keywords. Knowledge graph, meta-learning, dialogue generation

1. Introduction

Data-driven neural dialogue systems usually learn from a massive amount of conversational corpus using End-to-End [1-2], without combining hand-crafted rules or templates. However, Sequence-to-sequence (seq2seq) model tend to generate generic or incoherent responses. Recently, in order to generating high-quality responses, external knowledge is employed in open-domain dialogue systems, including unstructured texts [3] or structured knowledge representation [4-7].

Knowledge graphs can enhance capability of generating informative and diverse conversational responses. Because of the high human annotation cost, a limited number of triples suffer from information insufficiency for response generation. Nonetheless, the model capability of zero-shot adaptation to dynamic knowledge graph has rarely been considered. Entities or relations in dynamic knowledge graphs are temporal and evolve as a single time scale process [8].

In this paper, we propose an improved adversarial meta-learning algorithm [9] to facilitate knowledge aware dialogue generation. Adversarial meta-learning is presented based on Model-Agnostic Meta-Learning (MAML) [10]. The key idea of this article is

[1] Corresponding Author: Junpeng Bao, School of Computer Science and Technology, Xi'an Jiaotong University, China. E-mail: baojp@mail.xjtu.edu.cn.

considering dynamic entities and relations as adversarial samples, and fully utilizing knowledge graph-based dialog data to learn an initialization which adapt to new knowledge triples quickly. By combining Qadpt [8], a seq2seq neural conversation model with copy mechanism [11], we implement the ADML algorithm to learn an optimal initialization. We evaluate and show that our model outperforms the state-of-the-art baselines (Qadpt and TAware). The main contributions are as following:

1. Using meta-learning for knowledge dialogue tasks on a limited number of triples, learning meta-parameters effectively to adapt to knowledge-aware dialogue system.

2. Studying how to quickly train a dynamic knowledge graph-based dialogue model using a small dataset with both clean and adversarial samples.

2. Related Work

Knowledge Graph-based Conversations. Recently, there exist several models utilizing structured knowledge including factoid [4, 12] or commonsense knowledge [5-6] for generating informative responses. Researchers constructed several knowledge-aware datasets [8, 13]. [6] used knowledge graph embedding methods (e.g., TransE [14]) to encode each triple. However, these works are limited by incomplete knowledge graph.

Meta-Learning. Meta-learning or learning-to-learn aims at adapting quickly to new tasks with few steps and small datasets based on an optimal initialization. Recently, it has been applied on few-shot learning, such as machine translation [15], dialogue system [16-17], language generation [18], etc. There are three categories of meta-learning: Metric-based [19-22]: learning a metric space. Policy-based [23-25]: learning a policy to update model parameters. 3. Optimization-based [26]: learning a model parameter initialization adapting quickly to new tasks.

Adversary Attack. An adversarial sample refers to an instance with perturbations that cause a model to make a false prediction. Currently, there are several studies of adversarial attacks [27-32]. [27] mainly explored the principle of adversarial sample attack. [29] introduced the method of generating adversarial samples and adversarial training.

3. KgDg-ADML: dynamic KG-based dialogue model with ADML

3.1. Problem Formulation

For knowledge graph-based dialogue model Mθ, the context X and response Y are paired with knowledge graph $\mathcal{G}$. The model Mθ is expected to generate a sentence that is not only similar to the ground-turth Y, but is consistent to the entities and relationships.The idea of ADML is to utilize tasks $\{\mathcal{T}1,\ldots,\mathcal{T}K\}$ and learn the model initialization adaptive to new task, and each task has a loss function Li and contains a dataset Di (D = {(xn, yn, $\mathcal{G}$), n = 1…N}) that is further splited into clean and adversarial samples as Dtrain clean i, Dtrain adv i, Dtest clean i, Dtest adv i. Then, we compute loss and perform gradient descent to find the optimal parameter as θ'clean i and θ'adv i respectively. In meta-update stage, we next find the optimal parameter θ depended on θ'.

3.2. Knowledge Graph-based Dialogue Model

In knowledge graph-based dialogue system, $\mathcal{G}$ = {H, R, T} refers to a knowledge graph, where H, T $\in \mathcal{V}$ (the set of entities), R is a set of relationships, (h, r, t) is a triplets. Given a message X = {x_1, x_2,…, x_m} and $\mathcal{G}$, the goal is to generate a sequence Y = {y_1, y_2, …, y_n}. The system consists of two stages: (1) knowledge selection: the model selects the entities $\mathcal{V}$ to maximize the following probability as candidates:

$$v_Y = \arg\max_v P\left(v \mid v_X, \mathrm{G}, X\right) \tag{1}$$

v_X refers to entitie retrieved from $\mathcal{G}$, which is connected to word in X. v_Y refers to the vertex; (2) knowledge aware dialogue generation: it estimates the probability:

$$P\left(Y \mid X, v_Y\right) = \prod_{t=1}^{n} P\left(y_t \mid y_{<t}, X, v_Y\right) \tag{2}$$

The Qadpt model [8] is constructed based on a seq2seq model incorporating knowledge reasoning. Given context x, the encoder output a vector e(x), the decoder decode a vector d_t based on the ground-truth or predicted y:

$$\mathrm{e}\left(x\right) = GRU\left(x_1 x_2 x_3 \ldots x_m\right) d_t = GRU\left(y_1 y_2 y_3 \ldots y_{t-1}, e\left(x\right)\right) \tag{3}$$

where d_t decides to copy knowledge graph entities or generic words, and generates the path matrix R_t for knowledge graph reasoning, as shown in the following formula. These two processes are considered as knowledge selection module.

$$\begin{aligned} p\left(\left\{KB, W\right\} \mid y_1 y_2 y_3 \ldots y_{t-1}, e\left(x\right)\right) \\ = soft\max\left(\phi\left(d_t\right)\right) \end{aligned} \tag{4}$$

$$\begin{aligned} w_t &= p\left(W \mid y_1 y_2 y_3 \ldots y_{t-1}, e\left(x\right)\right) \\ c_t &= p\left(KB \mid y_1 y_2 y_3 \ldots y_{t-1}, e\left(x\right)\right) \\ o_t &= \left\{c_t k_t; w_t\right\}, \end{aligned} \tag{5}$$

where the probability c_t refers to the controller which is used to choose entities V from knowledge graph, while the probability $1 - c_t$ is used to choose generic words W. φ is a fully connected neural network, and k_t is the predicted distribution over knowledge graph entities V, and o_t is the produced distribution over all vocabularies.

$$R_t = soft\max\left(\theta\left(d_t\right)\right) \tag{6}$$

$$A_{i,j,\gamma} = \begin{cases} 1, \left(h_i, r_j, t_\gamma\right) \in K \\ 0, \left(h_i, r_j, t_\gamma\right) \notin K \end{cases} \tag{7}$$

$$T_t = R_t A \tag{8}$$

where θ is a linear transformation operation, R_t refers to the probability distribution of each head h ∈ V choosing each relation type r ∈ L. T_t is a transition matrix. A is an adjacency matrix which is a binary matrix indicating if the relations between two entities exist.

$$k_t = s^T \left(T_t\right)^N \tag{9}$$

where s is a binary vector used to indicate whether each entity exists in the message x. s is multiplied by the transition matrix T_t to produce K_t which is a probability distribution over knowledge entities, where N refers to multi-hop reasoning.

3.3. Improved ADML for KG-based Dialogue Generation

ADML is able to learn the varying correlation between clean and adversarial samples to obtain a better and robust initialization. The ADML is shown in Figure 1. For each task $\mathcal{T}_i$, in the inner gradient update process, ADML updates θ' to the direction of the adversarial subspace (purple color) as well as clean subspace (red color) to reach two points $\theta'_{adv\,i}$ and $\theta'_{clean\,i}$ respectively. Then in the meta-update stage, based on $\theta'_{adv\,i}$ and $\theta'_{clean\,i}$, ADML further optimizes θ' to reach the optimal point θ_i^*, which is expected to fall into the intersection of two subspaces.

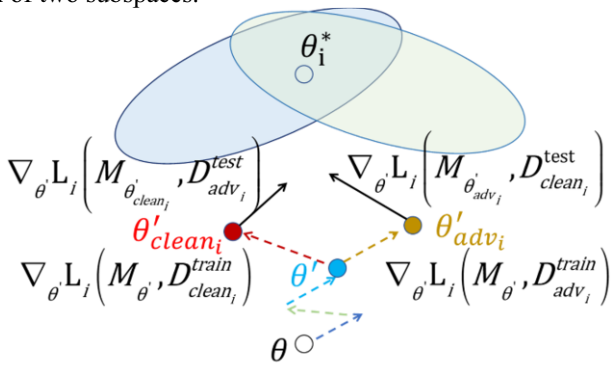

Figure 1. Illustration of design philosophy of improved ADML.

We denote the model Qadpt as M_θ parameterized by θ, which is updated iteratively. At each step, we sample a batch of tasks $\{\mathcal{T}_1, \ldots, \mathcal{T}_K\}$ containing support set D_i^{train} and query set D_i^{train} that are further splited into $D^{train}_{clean\,i}$, $D^{train}_{adv\,i}$. Then the model updates the parameters by k (k≥1) gradient descent steps for each task $\mathcal{T}_i$ in the following equations.

$$\theta_{clean_i}^{k} = \theta_{clean_i}^{k-1} - \alpha_1 \nabla_{\theta_{clean_i}^{k-1}} \mathrm{L}_{\mathrm{T}_i} \left(M_{\theta_{clean_i}^{k-1}}, D_{clearn_i}^{train} \right) \tag{10}$$

$$\theta_{adv_i}^{k+1} = \theta_{adv_i}^{k} - \alpha_2 \nabla_{\theta_{adv_i}^{k}} \mathrm{L}_{\mathrm{T}_i} \left(M_{\theta_{adv_i}^{k}}, D_{adv_i}^{train} \right) \tag{11}$$

$$\mathrm{L}_{\mathrm{T}_i} \left(\theta \right) = -\sum_{t=1}^{n} \log o_t \left(y_t \right) \tag{12}$$

where $\mathrm{L}_{\mathcal{T}\mathrm{i}}$ is the loss function for task $\mathcal{T}_\mathrm{i}$, α_1 and α_2 are the inner learning rate.

In the meta-update stage, we update the model parameters θ by optimizing the meta-objective function:

$$\min_{\theta} \sum_{\mathrm{T}_i \square \mathrm{T}} \mathrm{L}_i \left(M_{\theta_{clean_i}^{k}}, D'_{adv_i} \right) = \min_{\theta} \sum_{\mathrm{T}_i \square \mathrm{T}} \mathrm{L}_i \left(M_{\theta_{clean_i}^{k-1} - \alpha_1 \nabla_{\theta_{clean_i}^{k-1}} \mathrm{L}_{\mathrm{T}_i} \left(M_{\theta_{clean_i}^{k-1}}, D_{clearn_i}^{train} \right)}, D'_{adv_i} \right) \tag{13}$$

$$\min_{\theta} \sum_{\mathrm{T}_i \square \mathrm{T}} \mathrm{L}_i \left(M_{\theta_{adv_i}^{k+1}}, D'_{c_i} \right) = \min_{\theta} \sum_{\mathrm{T}_i \square \mathrm{T}} \mathrm{L}_i \left(M_{\theta_{adv_i}^{k} - \alpha_2 \nabla_{\theta_{adv_i}^{k}} \mathrm{L}_{\mathrm{T}_i} \left(M_{\theta_{adv_i}^{k}}, D_{adv_i}^{train} \right)}, D'_{c_i} \right) \tag{14}$$

The Meta-update of model M_θ is to update θ according to:

$$\theta = \theta - \beta_1 \nabla_{\theta} \sum_{\mathrm{T}_i \square p(\mathrm{T})} \mathrm{L}_{\mathrm{T}_i} \left(M_{\theta'_{clean_i}}, D_{clean_i}^{train} \right) \tag{15}$$

$$\theta = \theta - \beta_2 \nabla_{\theta} \sum_{\mathrm{T}_i \square p(\mathrm{T})} \mathrm{L}_{\mathrm{T}_i} \left(M_{\theta'_{adv_i}}, D_{adv_i}^{train} \right) \tag{16}$$

4. Experiments

4.1. Dataset and Evaluation Metrics

To compare with the state-of-the-art dialogue model, Qadpt [8] and TAware [11], we use the dataset HGZHZ, which first introduced in Qadpt. To verify whether our model can generate a more consistent and coherent response, there are five main metrics in our experiments including BLEU, PPL, DISTINCT1/2/3/4 to automatically evaluate the fluency, relevance, diversity, etc. The BLEU evaluates whether the generated response is also part of the task. PPL is a measurement of how well our model predicts a sample. DISTINCT measures the diversity of generated response.

4.2. Implementation Details

We set the learning rates as 0.001. In adversarial meta-learning stage, we set the num_task size as 4, support set size as 3, query set size as 4. We choose the one-layer GRU networks with a hidden size of 256 to construct the encoder and decoder. The model is optimized using Adam. We split data set to 4 buckets.

4.3. Results and Analysis

Table 1 summarizes the experimental results. We directly compare with the best results shown in [8]. We can observe that although TAware+multi method is overall better than the other models for KW/Generic, our model significantly outperforms other baselines. We also found that the methods using multi-hops reasoning technology outperform those without using multi-hops reasoning. It can be seen that our model has better capabilities on entities prediction.

Table 1. The results of entities prediction.

	KWAcc	KW/Generic		Generated-KW	
		Recall	Precision	Recall	Precision
TAware	50.21	44.40	**35.50**	49.18	76.72
+multi	57.71	**68.61**	28.70	44.50	90.70
Qadpt	57.61	38.24	28.31	44.50	90.70
+multi	57.40	51.97	28.43	44.50	**91.22**
Our model	**59.37**	40.37	**34.15**	**47.82**	**90.07**

Table 2. The results of responses generation with BLEU, perplexity (PPL), distinct scores (1-gram to 4-gram).

Models	BLEU	PPL	Dist1	Dist2	Dist3	Dist4
TAware	14.14	90.11	0.011	0.061	0.135	0.198
+multi	13.34	**80.48**	0.022	0.122	0.122	0.239
Qadpt	14.52	88.24	0.013	0.081	0.169	0.242
+multi	**15.47**	86.65	0.021	0.129	0.259	0.342
Our model	**14.95**	**82.49**	**0.031**	**0.157**	**0.312**	**0.415**

To evaluate the generated sentence quality, Table 2 presents the BLEU scores, perplexity (PPL) scores, and DISTINCT-N (**Dist**N) scores. The results show that our model can achieve a high consistency score, which is better than TAware, TAware+multi, Qadpt, and slightly less than Qadpt+multia. We can observe that our model significantly outperforms baselines in PPL. It can be seen that our method has significantly better performances. In summary, our model can better control the generation to maintain its coherence, fluency, relevance, and diversity with the dialog history and knowledge graphs.

As shown in Table 3, we only utilize very small datasets to compare our model with the Qadpt. We can see that our model achieves better results, which can prove that proposed method is robust to very small datasets.

Table 3. The results of responses generation with BLEU, PPL, distinct scores with very few datasets.

Samples	Method	PPL	BLEU	Dist1	Dist2	Dist3	Dist4
300	Qadpt	15240.94	19.35	0.266	0.716	0.842	0.827
	Our model	**1206.18**	**12.93**	**0.084**	**0.167**	**0.194**	**0.202**

5. Conclusion and Future

This paper proposes an algorithm for formulating dynamic knowledge graph as a problem of adversarial attack, focusing on the task of know- ledge aware dialogue generation. We use adversarial meta-gradients to find the optimal initialization that is robust to changed KG path and can adapt to very small datasets. We achieve baseline results on HGZHZ comparing to several state-of-the-art models. Experimental results show that our knowledge graph-based dialogue generation model can make full use of knowledge triples to generate informative response. Our model also provides promising potential extension, such as applying and data. We also plan to combine structural and non-structural knowledge to generate more content rich responses.

Acknowledgements

The work is supported by the CAAI-Huawei Mind-spore Fund.

References

[1] Sutskever I, Vinyals O, Le QV. Sequence to sequence learning with neural networks. Advances in neural information processing systems. 2014; 27.
[2] Serban I, Sordoni A, Bengio Y, Courville A, Pineau J. Building end-to-end dialogue systems using generative hierarchical neural network models. In Proceedings of the AAAI conference on artificial intelligence. 2016 March; 30(1).
[3] Ghazvininejad M, Brockett C, Chang MW, Dolan B. Gao J, Yih WT, Galley M. A knowledge-grounded neural conversation model. In Proceedings of the AAAI Conference on Artificial Intelligence. 2018 April; 32(1).
[4] Liu S, Chen H, Ren Z, Feng Y, Liu Q, Yin D. Knowledge diffusion for neural dialogue generation. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics. 2018 July; 1:1489-1498.
[5] Young T, Cambria E, Chaturvedi I, Zhou H, Biswas S, Huang M. Augmenting end-to-end dialogue systems with commonsense knowledge. In Proceedings of the AAAI conference on artificial intelligence. 2018 April; 32(1).
[6] Zhou H, Young T, Huang M, Zhao H, Xu J, Zhu X. Commonsense knowledge aware conversation generation with graph attention. In IJCAI. 2018 July; 4623-4629.
[7] Liu Z, Niu Z. Y, Wu H, Wang H. Knowledge Aware Conversation Generation with Explainable Reasoning over Augmented Graphs. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP). 2019.
[8] Tuan YL, Chen YN, Lee HY. DyKgChat: Benchmarking Dialogue Generation Grounding on Dynamic Knowledge Graphs. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP). 2019.
[9] Yin C, Tang J, Xu Z, Wang Y. Adversarial Meta-Learning. arXiv preprint arXiv:1806.03316. 2018.
[10] Finn C, Abbeel P, Levine S. Model-agnostic meta-learning for fast adaptation of deep networks. In International conference on machine learning. 2017 July; 1126-1135.
[11] Xing C, Wu W, Wu Y, Liu J, Huang Y, Zhou M, Ma WY. Topic aware neural response generation. In Proceedings of the AAAI conference on artificial intelligence. 2017 February; 31(1).
[12] Xu Z, Liu B, Wang B, Sun C, Wang X. Incorporating loose-structured knowledge into lstm with recall gate for conversation modeling. In Proceedings of IJCNN. 2017; 3506–3513.
[13] Wu W, Guo Z, Zhou X, Wu H, Zhang X, Lian R, Wang H. Proactive Human-Machine Conversation with Explicit Conversation Goal. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics. 2019; 3794-3804.
[14] Bordes A, Usunier N, Garcia-Duran A, Weston J, Yakhnenko O. Translating embeddings for modeling multi-relational data. Advances in neural information processing systems. 2013; 26.
[15] Gu J, Wang Y, Chen Y, Cho K, Li VO. Meta-learning for low-resource neural machine translation. arXiv preprint arXiv:1808.08437. 2018.

[16] Qian K, Yu Z. Domain Adaptive Dialog Generation via Meta Learning. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics. 2019; 2639–2649.
[17] Madotto A, Lin Z, Wu CS, Fung P. Personalizing dialogue agents via meta-learning. In Proceedings of the 57th annual meeting of the association for computational linguistics. 2019 July; 5454-5459.
[18] Huang PS, Wang C, Singh R, Yih WT, He X. Natural language to structured query generation via meta-learning. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Association for Computational Linguistics. 2018; 2: 732–738.
[19] Koch G, Zemel R, Salakhutdinov R. Siamese neural networks for one-shot image recognition. In ICML deep learning workshop. 2015 July; 2(1).
[20] Vinyals O, Blundell C, Lillicrap T, Wierstra D. Matching networks for one shot learning. Advances in neural information processing systems. 2016; 29:3630–3638.
[21] Santoro A, Bartunov S, Botvinick M, Wierstra D, Lillicrap T. Meta-learning with memory-augmented neural networks. In International conference on machine learning. 2016 June; 1842-1850.
[22] Sung F, Yang Y, Zhang L, Xiang T, Torr PH, Hospedales TM. Learning to compare: Relation network for few-shot learning. In Proceedings of the IEEE conference on computer vision and pattern recognition. 2018; 1199-1208.
[23] Andrychowicz M, Denil M, Gomez S, Hoffman MW, Pfau D, Schaul T, De Freitas N. Learning to learn by gradient descent by gradient descent. Advances in neural information processing systems. 2016; 29.
[24] Munkhdalai T, Yu H. Meta networks. In International Conference on Machine Learning. 2017; 2554–2563.
[25] Mishra N, Rohaninejad M, Chen X, Abbeel P. Meta-learning with temporal convolutions. arXiv preprint arXiv:1707.03141. 2017.
[26] Yoon J, Kim T, Dia O, Kim S, Bengio Y, Ahn S. Bayesian model-agnostic meta-learning. Advances in Neural Information Processing Systems. 2018; 7342–7352.
[27] Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. In International Conference on Learning Representations (ICLR). 2015.
[28] Kurakin A, Goodfellow I, Bengio S. Adversarial machine learning at scale. In International Conference on Learning Representations (ICLR). 2017.
[29] Kurakin A, Goodfellow I, Bengio S. Adversarial examples in the physical world. In International Conference on Learning Representations (ICLR). 2017.
[30] Ilyas A, Santurkar S, Tsipras D, Engstrom L, Tran B, Madry A. Adversarial examples are not bugs, they are features. Advances in neural information processing systems. 2019; 32.
[31] Yuan X, He P, Zhu Q, Li X. Adversarial examples: Attacks and defenses for deep learning. IEEE transactions on neural networks and learning systems. 2019; 30(9), 2805-2824.
[32] Zügner D, Akbarnejad A, Günnemann S. Adversarial attacks on neural networks for graph data. In Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining. 2018, July; 2847-2856.