

An Auto-Encoder Matching Model for Learning Utterance-Level Semantic Dependency in Dialogue Generation

Liangchen Luo*, Jingjing Xu*, Junyang Lin, Qi Zeng, Xu Sun

MOE Key Lab of Computational Linguistics, School of EECS, Peking University
{luolc, jingjingxu, linjunyang, pkuzengqi, xusun}@pku.edu.cn

Abstract

Generating semantically coherent responses is still a major challenge in dialogue generation. Different from conventional text generation tasks, the mapping between inputs and responses in conversations is more complicated, which highly demands the understanding of utterance-level semantic dependency, a relation between the whole meanings of inputs and outputs. To address this problem, we propose an Auto-Encoder Matching (AEM) model to learn such dependency. The model contains two auto-encoders and one mapping module. The auto-encoders learn the semantic representations of inputs and responses, and the mapping module learns to connect the utterance-level representations. Experimental results from automatic and human evaluations demonstrate that our model is capable of generating responses of high coherence and fluency compared to baseline models.¹

1 Introduction

Automatic dialogue generation task is of great importance to many applications, ranging from open-domain chatbots (Higashinaka et al., 2014; Vinyals and Le, 2015; Li et al., 2016, 2017a; Su et al., 2018) to goal-oriented technical support agents (Bordes and Weston, 2016; Zhou et al., 2017; Asri et al., 2017). Recently there is an increasing amount of studies about purely data-driven dialogue models, which learn from large corpora of human conversations without hand-crafted rules or templates. Most of them are based on the sequence-to-sequence (Seq2Seq) framework (Sutskever et al., 2014) that maximizes the probability of gold responses given the previous dialogue turn. Although such methods offer great

promise for generating fluent responses, they still suffer from the poor semantic relevance between inputs and responses (Xu et al., 2018). For example, given “What’s your name” as the input, the models generate “I like it” as the output.

Recently, the neural attention mechanism (Luong et al., 2015; Vaswani et al., 2017) has been proved successful in many tasks including neural machine translation (Ma et al., 2018b) and abstractive summarization (Lin et al., 2018), for its ability of capturing word-level dependency by associating a generated word with relevant words in the source-side context. Recent studies (Mei et al., 2017; Serban et al., 2017) have applied the attention mechanism to dialogue generation to improve the dialogue coherence. However, conversation generation is a much more complex and flexible task as there are less “word-to-words” relations between inputs and responses. For example, given “Try not to take on more than you can handle” as the input and “You are right” as the response, each response word can not find any aligned words from the input. In fact, this task requires the model to understand the utterance-level dependency, a relation between the whole meanings of inputs and outputs. Due to the lack of utterance-level semantic dependency, the conventional attention-based methods that simply capture the word-level dependency achieve less satisfying performance in generating high-quality responses.

To address this problem, we propose a novel **Auto-Encoder Matching** model to learn utterance-level dependency. First, motivated by Ma et al. (2018a), we use two auto-encoders to learn the semantic representations of inputs and responses in an unsupervised style. Second, given the utterance-level representations, the mapping module is taught to learn the utterance-level dependency. The advantage is that by explicitly separating representation learning and dependency

*Equal Contribution

¹The code is available at <https://github.com/lancopku/AMM>

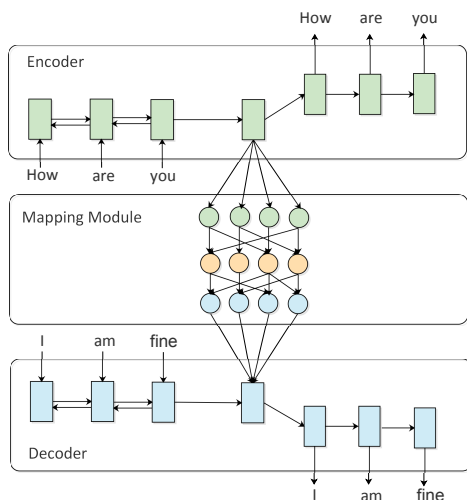


Figure 1: An illustration of the Auto-Encoder Matching model. The encoder and decoder are two auto-encoders that are responsible for learning the semantic representations. The mapping module is responsible for learning the utterance-level dependency.

learning, the model has a stronger modeling ability compared to traditional Seq2Seq models. Experimental results show that our model substantially outperforms baseline methods in generating high-quality responses.

Our contributions are listed as follows:

- To promote coherence in dialogue generation, we propose a novel Auto-Encoder Matching model to learn the utterance-level dependency.
- In our proposed model, we explicitly separate utterance representation learning and dependency learning for a better expressive ability.
- Experimental results on automatic evaluation and human evaluation show that our model can generate much more coherent text compared to baseline models.

2 Approach

In this section, we introduce our proposed model. An overview is presented in Section 2.1. The details of the modules are shown in Sections 2.2, 2.3 and 2.4. The training method is introduced in Section 2.5.

2.1 Overview

The proposed model contains three modules: an encoder, a decoder, and a mapping module, as shown in Figure 1.

In general, our model is different from the conventional sequence-to-sequence models. The encoder and decoder are both implemented as auto-encoders (Baldi, 2012). They learn the internal representations of inputs and target responses, respectively. In addition, a mapping module is built to map the internal representations of the input and the response.

2.2 Encoder

The encoder E_θ is an unsupervised auto-encoder based on Long Short Term Memory Networks (LSTM) (Hochreiter and Schmidhuber, 1996). As it is essentially a LSTM-based Seq2Seq model, we name the encoder and decoder of the auto-encoder “source-encoder” and “source-decoder”. To be specific, the encoder E_θ receives the source text $x = \{x_1, x_2, \dots, x_n\}$, and encodes it to an internal representation h , and then decodes h to a new sequence $\tilde{x} = \{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n\}$ for the reconstruction of the input. We extract the hidden state h as the semantic representation. The encoder E_θ is trained to reduce the reconstruction loss, whose loss function is defined as follows:

$$J_1(\theta) = -\log P(\tilde{x}|x; \theta) \quad (1)$$

where θ refers to the parameters of the encoder E_θ .

2.3 Decoder

Similar to the encoder, our decoder D_ϕ is also a LSTM-based auto-encoder. However, as there is no target text provided in the testing stage, we propose the customized implementation, which is illustrated in Section 2.5. Here in the introduction of the decoder, we do not provide the testing details. Similarly, we name the encoder and decoder of the auto-encoder “target-encoder” and “target-decoder”. The target-encoder receives the target $y = \{y_1, y_2, \dots, y_n\}$ and encodes it to a utterance-level semantic representation s , and then decodes s to a new sequence to approximate the target text. The loss function is identical to that of the encoder:

$$J_2(\phi) = -\log P(\tilde{y}|y; \phi) \quad (2)$$

2.4 Mapping Module

As our model is constructed for dialogue generation, we design the mapping module to ensure that the generated response is semantically consistent

with the source. There are many matching models that can be used to learn such dependency relations (Hu et al., 2014; Guo et al., 2016; Pang et al., 2016; Chen et al., 2018). For simplicity, we only use a simple feedforward network for implementation. The mapping module M_γ transforms the source semantic representation h to a new representation t . To be specific, we implement a multi-layer perceptron (MLP) $g(\cdot)$ for M_γ and train it by minimizing the L2-norm loss $J_3(\gamma)$ of the transformed representation t and the semantic representation of target response s :

$$\begin{aligned} t &= g(h) \\ J_3(\gamma) &= \frac{1}{2} \|t - s\|_2^2 \end{aligned} \quad (3)$$

2.5 Training and Testing

In the testing stage, given an input utterance, the encoder E_θ , the decoder D_ϕ , and the matching module M_γ work together to produce a dialogue response. The source-encoder first receives the input x and encodes it to a semantic representation h of the source utterance. Then, the mapping module transforms h to t , a target response representation. Finally, t is sent to the target-decoder for response generation.

In the training stage, besides the auto-encoder loss and the mapping loss, we also use an end-to-end loss $J_4(\theta, \phi, \gamma)$:

$$\begin{aligned} J_4(\theta, \phi, \gamma) &= -\log P(y|x; \theta, \phi, \gamma) \\ &= -\sum_{t=1}^T \log P(y_t|x, y_{1..t-1}; \theta, \phi, \gamma) \end{aligned} \quad (4)$$

$$(5)$$

where x is the source input, y is the target response, and T is the length of response sequence. The model learns to generate \tilde{y} to approximate y by minimizing the reconstruction losses $J_1(\theta)$ and $J_2(\phi)$, the mapping loss $J_3(\gamma)$, and the end-to-end loss $J_4(\theta, \phi, \gamma)$. The details are illustrated below:

$$\begin{aligned} J &= \lambda_1 [J_1(\theta) + J_2(\phi)] + \lambda_2 J_3(\gamma) \\ &\quad + \lambda_3 J_4(\theta, \phi, \gamma) \end{aligned} \quad (6)$$

where J refers to the total loss, and λ_1 , λ_2 , and λ_3 are hyperparameters.

3 Experiment

We conduct experiments on a high-quality dialogue dataset called **DailyDialog** built by Li et al.

(2017b). The dialogues in the dataset reflect our daily communication and cover various topics about our daily life. We split the dataset into three parts with 36.3K pairs for training, 11.1K pairs for validation, and 11.1K pairs for testing.

3.1 Experimental Details

For dialogue generation, we set the maximum length to 15 words for each generated sentence. Based on the performance on the validation set, we set the hidden size to 512, embedding size to 64 and vocabulary size to 40K for baseline models and the proposed model. The parameters are updated by the Adam algorithm (Kingma and Ba, 2014) and initialized by sampling from the uniform distribution $([-0.1, 0.1])$. The initial learning rate is 0.002 and the model is trained in mini-batches with a batch size of 256. λ_1 and λ_3 are set to 1 and λ_2 is set to 0.01 in Equation (6). It is important to note that for a fair comparison, we re-implement the baseline models with the best settings on the validation set. After fixing the hyperparameters, we combine the training and validation sets together as a larger training set to produce the final model.

3.2 Results

We use BLEU (Papineni et al., 2002), to compare the performance of different models, and use the widely-used BLEU-4 as our main BLEU score. The results are shown in Table 1. The proposed AEM model significantly outperforms the Seq2Seq model. It demonstrates the effectiveness of utterance-level dependency on improving the quality of generated text. Furthermore, we find that the utterance-level dependency also benefits the learning of word-level dependency. The improvement from the AEM model to the AEM+Attention model² is 0.68 BLEU-4 point. It is much more obvious than the improvement from the Seq2Seq model to the Seq2Seq+Attention, which is 0.29 BLEU-4 point.

We also report the diversity of the generated responses by calculating the number of distinct unigrams, bigrams, and trigrams. The results are shown in Table 2. We find that the AEM model achieves significant improvement on the diversity of generated text. The number of unique trigram of the AEM model is almost six times more

²With the additional attention mechanism, the outputs of attention-based decoder and our decoder are combined together to predict the probability of response words.

Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4
Seq2Seq	12.43	4.57	2.69	1.84
AEM	13.55	4.89	3.04	2.16
Seq2Seq+Attention	13.63	4.99	3.05	2.13
AEM+Attention	14.17	5.69	3.78	2.84

Table 1: BLEU scores for the AEM model and the Seq2Seq model.

Models	Dist-1	Dist-2	Dist-3
Seq2Seq	0.8K	2.7K	5.5K
AEM	3.1K	14.8K	31.2K
Seq2Seq+Attention	2.5K	13.6K	34.6K
AEM+Attention	3.3K	23.2K	53.9K

Table 2: Diversity scores for the AEM model and the Seq2Seq model. Dist-1, Dist-2, and Dist-3 are respectively the number of distinct unigrams, bigrams, and trigrams in the generated text.

Models	Fluency	Coherence	G-Score
Seq2Seq	6.97	3.51	4.95
AEM	8.11	4.18	5.82
Seq2Seq+Attention	5.11	3.30	4.10
AEM+Attention	7.92	4.97	6.27

Table 3: Human evaluation results of the AEM model and the Seq2Seq model.

than that of the Seq2Seq model. Also, it should be noticed that the attention mechanism performs almost the same compared to the AEM model (31.2K vs. 34.6K in terms of Dist-3), which indicates that the utterance-level dependency and the word-level dependency are both indispensable for dialogue generation. Therefore, by combining the two dependencies together, the AEM+Attention model achieves the best results. Such improvements are expected. With the increase of the relevance of the generated text, it gets harder for the model to generate repeated responses. In our experimental results, the number of repetitive “I don’t know” in the AEM+Attention model is reduced by 50% compared to the Seq2Seq model.

For dialogue generation, human evaluation is more convincing, so we also report human evaluation results on the test set. We randomly choose 100 utterances in daily communication style for the human evaluation, each of which is sent to different models to generate responses. The results are distributed to the annotators who have no knowledge about which model the sentence is from. All annotators have linguistic background. They are asked to score the generated responses in terms of fluency and coherence. Fluency rep-

Input 1: <i>My name is David. What is my name?</i>
Seq2Seq: <i>Ours, maria.</i>
AEM: <i>Pardon me? It is unreasonable!</i>
Input 2: <i>I’m going to kill you!</i>
Seq2Seq: <i>I know. I’m a good man.</i>
AEM: <i>What?!?!?! No!!</i>
Input 3: <i>I hate you.</i>
Seq2Seq: <i>I’m not.</i>
AEM: <i>You do? you’re a sweet girl.</i>
Input 4: <i>How are you?</i>
Seq2Seq: <i>Fine.</i>
AEM: <i>Fine. I’m good.</i>
Input 5: <i>Can you say ‘Hello’?</i>
Seq2Seq: <i>Yeah, I guess.</i>
AEM: <i>Sure, I guess.</i>

Table 4: Examples generated by the AEM model and the Seq2Seq model.

resents whether each sentence is in correct grammar. Coherence evaluates whether the generated response is relevant to the input. The score ranges from 1 to 10 (1 is very bad and 10 is very good). To evaluate the overall performance, we use the geometric mean of fluency and coherence as the final evaluation metric.

Table 3 shows the results of human evaluation. The inter-annotator agreement is satisfactory considering the difficulty of human evaluation. The Pearson’s correlation coefficient is 0.69 on coherence and 0.57 on fluency, with $p < 0.0001$. First, it is clear that the AEM model outperforms the Seq2Seq model with a large margin, which proves the effectiveness of the AEM model on

generating high quality responses. Second, it is interesting to note that with the attention mechanism, the coherence is decreased slightly in the Seq2Seq model but increased significantly in the AEM model. It suggests that the utterance-level dependency greatly benefits the learning of word-level dependency. Therefore, it is expected that the AEM+Attention model achieves the best G-score.

Table 4 shows the examples generated by the AEM model and the Seq2Seq model. For easy questions (ex. 4 and ex. 5), they both perform well. For hard questions (ex. 1 and ex. 2), the proposed model obviously outperforms the Seq2Seq model. It shows that the utterance-level dependency learned by the proposed model is useful for handling complex inputs.

3.3 Error Analysis

Although our model achieves the best performance, there are still several failure cases. We find that the model performs badly for the inputs with unseen words. For instance, given “Bonjour” as the input, it generates “Stay out of here” as the output. It shows that the proposed model is sensitive to the unseen utterance representations. Therefore, we would like to explore more approaches to address this problem in the future work. For example, the auto-encoders can be replaced by variational auto-encoders to ensure that the distribution of utterance representations is normal, which has a better generalization ability.

4 Conclusion

In this work, we propose an Auto-Encoder Matching model to learn the utterance-level semantic dependency, a critical dependency relation for generating coherent and fluent responses. The model contains two auto-encoders that learn the utterance representations in an unsupervised way, and a mapping module that builds the mapping between the input representation and response representation. Experimental results show that the proposed model significantly improves the quality of generated responses according to automatic evaluation and human evaluation, especially in coherence.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China (No. 61673028). We thank all reviewers for providing the construc-

tive suggestions. Xu Sun is the corresponding author of this paper.

References

- Layla El Asri, Hannes Schulz, Shikhar Sharma, Jeremie Zumer, Justin Harris, Emery Fine, Rahul Mehrotra, and Kaheer Suleman. 2017. Frames: a corpus for adding memory to goal-oriented dialogue systems. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue, Saarbrücken, Germany, August 15-17, 2017*, pages 207–219.
- Pierre Baldi. 2012. Autoencoders, unsupervised learning, and deep architectures. In *Unsupervised and Transfer Learning - Workshop held at ICML 2011, Bellevue, Washington, USA, July 2, 2011*, pages 37–50.
- Antoine Bordes and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *CoRR*, abs/1605.07683.
- Deli Chen, Shuming Ma, Pengcheng Yang, and Xu Sun. 2018. Identifying high-quality chinese news comments based on multi-target text matching model. *CoRR*, abs/1808.07191.
- Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*, pages 55–64.
- Ryuichiro Higashinaka, Kenji Imamura, Toyomi Meguro, Chiaki Miyazaki, Nozomi Kobayashi, Hiroaki Sugiyama, Toru Hirano, Toshiro Makino, and Yoshihiro Matsuo. 2014. Towards an open-domain conversational system fully based on natural language processing. In *COLING 2014, 25th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, August 23-29, 2014, Dublin, Ireland*, pages 928–939.
- Sepp Hochreiter and Jürgen Schmidhuber. 1996. LSTM can solve hard long time lag problems. In *NIPS, 1996*, pages 473–479.
- Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2042–2050.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016*,

- Austin, Texas, USA, November 1-4, 2016, pages 1192–1202.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017a. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2157–2169.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017 - Volume 1: Long Papers*, pages 986–995.
- Junyang Lin, Xu Sun, Shuming Ma, and Qi Su. 2018. Global encoding for abstractive summarization. *CoRR*, abs/1805.03989.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1412–1421.
- Shuming Ma, Xu Sun, Junyang Lin, and Houfeng Wang. 2018a. Autoencoder as assistant supervisor: Improving text representation for chinese social media text summarization. *CoRR*, abs/1805.04869.
- Shuming Ma, Xu Sun, Yizhong Wang, and Junyang Lin. 2018b. Bag-of-words as target for neural machine translation. *CoRR*, abs/1805.04871.
- Hongyuan Mei, Mohit Bansal, and Matthew R. Walter. 2017. Coherent dialogue with attention-based language models. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3252–3258.
- Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2793–2799.
- 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.
- Iulian Vlad Serban, Tim Klinger, Gerald Tesauro, Karthik Talamadupula, Bowen Zhou, Yoshua Bengio, and Aaron C. Courville. 2017. Multiresolution recurrent neural networks: An application to dialogue response generation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3288–3294.
- Hui Su, Xiaoyu Shen, Pengwei Hu, Wenjie Li, and Yun Chen. 2018. Dialogue generation with GAN. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- 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 6000–6010.
- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. *CoRR*, abs/1506.05869.
- Jingjing Xu, Xuancheng Ren, Junyang Lin, and Xu Sun. 2018. Diversity-promoting gan: A cross-entropy based generative adversarial network for diversified text generation. In *EMNLP, 2018*.
- Li Zhou, Kevin Small, Oleg Rokhlenko, and Charles Elkan. 2017. End-to-end offline goal-oriented dialog policy learning via policy gradient. *CoRR*, abs/1712.02838.