

Know More about Each Other: Evolving Dialogue Strategy via Compound Assessment

Siqi Bao, Huang He, Fan Wang, Rongzhong Lian and Hua Wu
Baidu Inc., China

{baosiqi, hehuang, wangfan04, lianrongzhong, wu_hua}@baidu.com

Abstract

In this paper, a novel Generation-Evaluation framework is developed for multi-turn conversations with the objective of letting both participants know more about each other. For the sake of rational knowledge utilization and coherent conversation flow, a dialogue strategy which controls knowledge selection is instantiated and continuously adapted via reinforcement learning. Under the deployed strategy, knowledge grounded conversations are conducted with two dialogue agents. The generated dialogues are comprehensively evaluated on aspects like informativeness and coherence, which are aligned with our objective and human instinct. These assessments are integrated as a compound reward to guide the evolution of dialogue strategy via policy gradient. Comprehensive experiments have been carried out on the publicly available dataset, demonstrating that the proposed method outperforms the other state-of-the-art approaches significantly.

1 Introduction

Intelligent dialogue systems have become popular in our daily life, such as the chat XiaoIce and the task-oriented Echo. These systems serve as smart agents to facilitate more effective interaction with users in various situations, like ticket booking or recreation offering. Primary dialogue systems (Vinyals and Le, 2015; Shang et al., 2015) try to mimic human beings to generate fluent utterances, whereas paying little attention to the intrinsic factors of human conversations: *exchanging information* and *enhancing interaction* (Li et al., 2017). Therefore, they are prone to generate dull and generic responses.

To address this problem, in recent years, several approaches have been developed to generate informative responses based on external knowledge. Recently, a knowledge grounded model is

proposed in Ghazvininejad et al. (2018), where relevant factual texts are encoded into memory and replies are decoded via attention mechanism. Instead of using unstructured text knowledge, CCM (Zhou et al., 2018) relies on structured knowledge to generate rich-information response. However, all these approaches are designed for the single-round settings. While applied to the real-world scenarios (where dialogues are conducted for multiple rounds), the dialogue quality will be severely limited due to the lack of coordination among different rounds.

As discussed above, one of the ultimate goals in human conversation is that information can be exchanged effectively through interaction. Particularly, we argue that successful multi-turn dialogues are determined by the joint experience of both participants in the conversation, i.e., both participants need to get aware of their counterparts and express themselves effectively. To this end, we propose the objective of letting both sides know more about each other. With this objective, a novel Generation-Evaluation framework is introduced for the multi-turn dialogues.

As the name Generation-Evaluation indicates, there are two fundamental modules in our framework. In the module of dialogue generation, a two-stage generative model is employed, where the dialogue strategy determines which knowledge to use for the current turn and the decoder uses this knowledge to produce the response. In the module of evaluation, the generated dialogues are assessed from the following two aspects: *informativeness*, which measures the effectiveness of information exchange and *coherence*, which reflects the response's suitability. Both modules are assembled within a unified reinforcement learning pipeline. The generation module simulates knowledge grounded conversations with two dialogue agents and receives compound reward from the

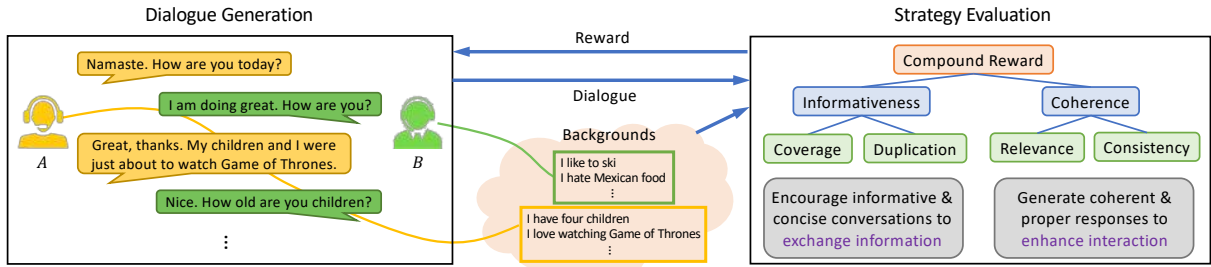


Figure 1: Framework overview. Left: dialogue generation. Right: strategy evaluation.

evaluation module. By keeping adapted for higher evaluation rewards, the generation module will be continuously evolving for better dialogue quality. As suggested in Yarats and Lewis (2018), applying reinforcement learning on the decoder might bring in adverse impacts on the linguistic quality. As such, in the generation module, the decoder is pre-trained with supervised learning and the dialogue strategy keeps evolving with reinforcement learning.

The contributions of this work are summarized as follows:

- With the objective of letting both participants know more about each other, we propose a novel Generation-Evaluation framework, which facilitates the generation of informative and coherent dialogues.
- To evaluate the effectiveness of dialogue strategy, two metrics are specially designed on informativeness and coherence, which are further integrated as a compound reward. Towards maximizing this reward, the strategy of knowledge selection is able to evolve via reinforcement learning.
- Intensive and extensive experiments have been carried out on PersonaChat. As compared with other state-of-the-art approaches, our method obtains superior performances on both automatic and human evaluations.

2 Methodology

2.1 Framework Overview

Our Generation-Evaluation framework is illustrated in Figure 1. Under the deployed strategy of knowledge selection, two dialogue agents introduce themselves alternately in accordance with corresponding backgrounds and make responses




Figure 2: Architecture of dialogue generation.

to their counterparts in a proper way. The generated dialogues together with the agents' backgrounds are collected for strategy evaluation in terms of two essential aspects: informativeness and coherence. Then these assessments are integrated as a compound reward, acting as the reinforcing signal for the evolution of knowledge interaction strategy.

In the following parts, we will first introduce the process of dialogue generation, present the metrics utilized in strategy evaluation and then describe the strategy evolution via compound assessment.

2.2 Dialogue Generation

The detailed network architecture of dialogue generation is illustrated in Figure 2. With the context and background knowledge as input, our dialogue strategy selects one piece of appropriate knowledge to generate informative and coherent response. The background $Z = \{z_1, z_2, \dots, z_M\}$ includes a set of knowledge, where a piece of knowledge z_i is presented by one sentence, such as "i like to ski". Utterance u_{t-1} is the last response from the other participant and the context $c_t = \text{concat}(u_1, u_2, \dots, u_{t-1})$ is the current conversation history.

It is worth noting that in our dialogue generation, the input context c_t is separated into two parts, with independent encoders employed for utterance u_{t-1} and context c_{t-1} respectively. The motivation to do so lies in two aspects: for the

sake of coherence, the knowledge utilized in t -th turn is supposed to be semantically related to the partner’s last utterance u_{t-1} ; to avoid repetition, the knowledge utilized in t -th turn should be dissimilar with the former dialogue history c_{t-1} .

After passing through the embedding layer and the encoders of gated recurrent unit (GRU) (Cho et al., 2014), the inputs obtain their corresponding feature representation: knowledge z_i^G , utterance u_{t-1}^G and context c_{t-1}^G . $\mathcal{Z}^G = \{z_1^G, z_2^G, \dots, z_M^G\}$ is the set of knowledge representation. With discriminative representations u_{t-1}^G , c_{t-1}^G and \mathcal{Z}^G obtained, the prior distribution over knowledge $p(\mathcal{Z}|c_t)$ can be estimated through MLP attention (MLP-ATT) (Bahdanau et al., 2015):

$$\begin{aligned} p(\mathcal{Z}|c_t) &= p(\mathcal{Z}|u_{t-1}) * 0.5 + p(\mathcal{Z}|c_{t-1}) * 0.5, \\ p(z_i|u_{t-1}) &= \text{softmax}(\text{MLP-ATT}(u_{t-1}^G, z_i^G)), \\ p(z_i|c_{t-1}) &= \text{softmax}(\text{MLP-ATT}(c_{t-1}^G, z_i^G)), \end{aligned} \quad (1)$$

where softmax is defined as $\text{softmax}(s_i) = e^{s_i} / \sum_j e^{s_j}$ (Sukhbaatar et al., 2015). And the computation of MLP-ATT is given as follows:

$$\text{MLP-ATT}(x, y) = V_1^T \tanh(xW_1 + yW_2),$$

where $W_1, W_2 \in \mathbb{R}^{d \times d}$ and $V_1 \in \mathbb{R}^d$ are the weight matrices. $p(\mathcal{Z}|c_t)$ is the probability distribution for knowledge selection and $\sum_{i=1}^M p(z_i|c_t) = 1$. (If $p(z_i|c_t) = 0.2$, it means that the probability to select knowledge z_i is 0.2.) According to the estimated prior probability distribution $p(\mathcal{Z}|c_t)$, one piece of knowledge can be sampled $z_i \sim p(\mathcal{Z}|c_t)$ and sent to the decoder for response generation $p(u_t|z_i, u_{t-1})$.

It is obvious that the key component for informative and coherent conversation is the appropriate knowledge selection, shown as Blue areas in Figure 2. Nevertheless, a high-fidelity decoder $p(u_t|z_i, u_{t-1})$, which is able to express the given knowledge accurately, is also indispensable. To this end, the pre-training is carried out using those target responses associated with ground-truth knowledge via supervised learning. The training data is in the format of $\{u_{t-1}, z_i, u_t\}$, where u_{t-1} is the last utterance from the partner, u_t is the target response and z_i is the ground truth knowledge used in u_t . Major steps in the pre-training are listed as follows: (1) the encoders convert the knowledge and utterance into z_i^G and u_{t-1}^G ; (2) the decoder tries to generate the response u_t based on the ground-truth knowledge z_i and last




Figure 3: Toy example of informativeness assessment: activation a_t records whether a piece of knowledge is expressed in u_t , coverage v_t keeps track of expressed knowledge and repetition d_t detects reiteration.

utterance u_{t-1} ; (3) parameters in the encoders and decoder (Gray areas) are optimized via supervised learning, with the loss functions defined in Zhao et al. (2017). For the rest of the parameters related to the knowledge selection strategy (Blue areas), they will keep evolving through Generation-Evaluation reinforcement learning, which will be discussed in detail.

2.3 Strategy Evaluation

Multi-turn knowledge grounded conversations are generated by two dialogue agents. To evaluate the effectiveness of deployed strategy, generated conversations and agents’ background knowledge are collected for evaluation and two metrics are judiciously designed – informativeness and coherence.

2.3.1 Informativeness

Information is a crucial ingredient in generating meaningful conversations. Although many approaches have been introduced to boost the generation of informative utterances, due to a lack of thorough control on effective information utilization, they are prone to generating repetitive utterances in multi-turn conversations. In this paper, we design a novel informativeness metric to measure the effective exploitation of information in the conversation level, which encourages extensive coverage and avoids unnecessary repetition.

To illustrate the informativeness assessment, a toy example is given in Figure 3. Assume that there are five pieces of background knowledge z_i within the conversation participants. For each generated utterance u_t , it will be assessed whether z_i is expressed by u_t or not, which can be approximately inferred through keyword matching (in the form of binary variable 0/1). Such estimation over the background knowledge is stored in the activation vector a_t . If relying on a_t as the informa-

tiveness metric, it is able to boost informative response generation on the utterance level. However, it inevitably produces repetitive responses due to the lack of information utilization control on the conversation level.

Inspired by the coverage mechanism in machine translation (Tu et al., 2016) and text summarization (See et al., 2017), we propose to maintain one coverage vector v_t to keep track of the activation on each piece of information during the conversation flow. From the toy example, it can be observed that the coverage vector v_t increases with the amount of expressed knowledge. In other words, a higher mean value of v_t indicates that the participants have expressed more background knowledge, which gives a better chance for them to know more about each other.

Although the coverage mechanism stimulates extensive knowledge expression, it still lacks effective and explicit control on the reiteration. For the sake of user experience, we also maintain one repetition vector d_t to detect information redundancy, whose estimation is carried out by jointly considering current information activation and last-step coverage status:

$$d_t = \min(a_t, v_{t-1}), \quad (2)$$

where the function $\min(\cdot)$ calculates the element-wise minimum value between two vectors. As shown in Figure 3, when utterance u_3 reiterates the same information as before, it does not increase knowledge coverage and leads to unnecessary repetition.

In summary, instead of focusing on the information activation of the single-round response, our informativeness metric considers the effective information utilization in the scope of multi-turn conversation. For a conversation with T turns, its informativeness is estimated as follows:

$$r_I = \text{mean}(v_T) - \sum_{t=1}^T \text{mean}(d_t), \quad (3)$$

where the function $\text{mean}(\cdot)$ calculates the mean value of a vector. By maintaining information coverage and internal repetition simultaneously, the conversation level informativeness is able to encourage informative and concise conversations.

2.3.2 Coherence

For the sake of natural interaction, coherence is another indispensable ingredient in strategy evaluation. In addition to relevance with the context,




Figure 4: Illustration of coherence assessment, where H-GRU refers to hierarchical GRU and the symbol \oplus denotes vector concatenation.

the coherence assessment also evaluates the conversation consistency with the backgrounds. The motivation to enforce background consistency is to confine the massive and loose interactive responses into a reasonable space. Considering that the essence of coherence is semantic relevance between two inputs and many deep learning based approaches have demonstrated their superiority at capturing semantic relevance, such as DSSM (Huang et al., 2013), SMN (Wu et al., 2017) and BERT (Devlin et al., 2018), we use a symmetric neural network for the coherence assessment in this paper.

As shown in Figure 4, for a generated utterance u_t , its coherence with the context c_t and corresponding backgrounds Z can be estimated through this symmetric network. The utterance is fed into the embedding layer, followed by gated recurrent unit (GRU) (Cho et al., 2014) and multilayer perceptron (MLP) to capture discriminative representation. As for the context and backgrounds, they are fed into the embedding layer and the hierarchical GRU for better feature extractions (Sordani et al., 2015), which are further concatenated together to obtain comprehensive representation. The final coherence is estimated as the inner product between two vectors:

$$r_{C_t} = \sigma(\text{MLP}(u_t^G) \cdot \text{MLP}([c_t^H, z^H])), \quad (4)$$

where $\text{MLP}(x) = \sigma(xW_1 + b_1)W_2 + b_2$.

$\sigma(\cdot)$ is the sigmoid activation, $[\cdot, \cdot]$ denotes vector concatenation and MLP includes two linear transformations with a sigmoid activation in between.

The above equation evaluates the coherence for each generated utterance u_t , by considering existing conversation history and corresponding background, which is further summed up over all utterances as conversation-level coherence assessment.

2.3.3 Compound Assessment

To provide a united reinforcement signal for strategy evolution, the informativeness and coherence assessments are further integrated as a compound reward. For a conversation τ with T turns, the compound assessment is defined as:

$$R(\tau) = \sum_{t=1}^T r_{C_t} + r_I. \quad (5)$$

The two intrinsic factors in human conversations – exchanging information and enhancing interaction have been included in our compound reward.

2.4 Strategy Evolution

From the perspective of reinforcement learning, the knowledge selection within a conversation can be regarded as sequential actions taken within a trajectory. As such, the objective of knowledge grounded dialogue generation can be written as:

$$\max J(\theta) = \mathbb{E}_{\tau \sim p(\tau; \theta)} R(\tau), \quad (6)$$

where θ refers to the network parameters of dialogue generation, $\tau \sim p(\tau; \theta)$ is a multi-turn conversation generated under the deployed strategy and $R(\tau)$ is the compound assessment of strategy evaluation. Gradient update of the above objective can be further derived as follows:

$$\begin{aligned} \nabla_{\theta} J(\theta) &= \sum_{t=1}^T \nabla_{\theta} \log(p(z_i|c_t)p(u_t|z_i, u_{t-1}))(R(\tau) - b), \\ &= \sum_{t=1}^T \nabla_{\theta} \log p(z_i|c_t)(R(\tau) - b) \\ &\quad + \sum_{t=1}^T \nabla_{\theta} \log p(u_t|z_i, u_{t-1})(R(\tau) - b), \end{aligned} \quad (7)$$

where b is the reward baseline estimated with K times Monte Carlo sampling: $b = \sum_k R(\tau^{(k)})/K$. In Equation (7), the first term is about the dialogue strategy of appropriate knowledge selection and the second term is about the decoding process with the selected knowledge. As suggested in (Lewis et al., 2017; Yarats and Lewis, 2018), applying reinforcement learning on the decoder might lead to poor linguistic quality. As such, in this paper, the focus is on the strategy evolution and gradient update is further simplified:

$$\nabla_{\theta} J(\theta) = \sum_{t=1}^T \nabla_{\theta} \log p(z_i|c_t)(R(\tau) - b). \quad (8)$$

The physical meaning of the above equation is given as follows: the strategies that lead to higher conversation rewards will be encouraged and those that result in lower conversation rewards will be suppressed.

As demonstrated in Equation (8), the network parameters related to dialogue strategy (Blue areas in Figure 2) will keep evolving via compound assessment. For the rest parameters, they are pre-trained with supervised learning and will be kept fixed during strategy evolution.

3 Experiments

3.1 Settings

All experiments have been carried out on the publicly available dataset – PersonaChat (Zhang et al., 2018), which provides both human annotated conversations and the participants’ background knowledge (persona profiles). PersonaChat has separated training and testing set. In total, there are 8,939 dialogues (131,438 turns) in the training set and 968 dialogues (15,024 turns) in the testing set. Comprehensive comparisons have been made to the following methods:

- Sequence to sequence with attention (Seq2Seq) (Vinyals and Le, 2015) is the classic response generation approach, without using any extra knowledge.
- The knowledge grounded memory network (Mem-Net) (Ghazvininejad et al., 2018) encodes text knowledge into memory to boost the generation of informative responses.
- The KG-Net (Lian et al., 2019) makes use of posterior knowledge distribution in the training process for accurate informative response generation and achieves the state-of-the-art results on PersonaChat.
- Li et al. (2016b) first employed reinforcement learning for dialogue generation (RL-DG), where simple Seq2Seq was used as the generation model. In the experiments, to improve RL-DG’s performance, KG-Net is utilized as the base model for informative generation.

In our strategic knowledge interaction, the parameters of knowledge encoder, utterance encoder and decoder were pre-trained with supervised learning. For the learnable parameters (Blue

areas in Figure 2), the context encoder was initialized with the utterance encoder and random initialization was employed for the rest layers¹. The training process was carried out using Adam optimizer, with a learning rate of $2e-4$. The conversation turns T was set to 8, batch size was set to 8 and Monte Carlo sampling times K was set to 16.

3.2 Experimental Results

The training curves of reinforcement learning are shown in Figure 5, which are the results averaged over 5 random seeds. The horizontal axis refers to the number of trained dialogues. The vertical axis stands for the compound episode reward, informativeness and coherence, respectively. These results demonstrate that all rewards increase stably within the training process and remarkable increments are achieved after convergence.

3.2.1 Automatic Evaluation

The experimental results with automatic measurements are summarized in Table 1, with highest value written in bold. Distinct-1/2 (Li et al., 2016a) measures the diversity of generated conversations, which is defined as the amount of distinct unigrams or bigrams divided by the total number of generated words. Knowledge-Recall/Precision/F1 (Dinan et al., 2019b) measures the informativeness of generated conversations with regarding to background knowledge, defined as:

$$\begin{aligned} \text{Recall} &= \frac{|W_G \cap W_K|}{|W_K|}, \\ \text{Precision} &= \frac{|W_G \cap W_K|}{|W_G|}, \\ \text{F1} &= 2 \times \frac{\text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}}, \end{aligned} \quad (9)$$

where W_G and W_K refer to the set of non-stop words in generated conversations and background knowledge.

From Table 1, it demonstrates that the proposed method obtains the best results. The distinct measurement indicates that more diverse words or phrases are produced by our method. The knowledge measurement verifies the effectiveness of our approaches on the knowledge utilization in multi-turn conversations. As compared with the state-of-the-art KG-Net, the knowledge F1 of our method

¹Our code and model will be released at <https://github.com/PaddlePaddle/models/tree/develop/PaddleNLP/Research/ACL2019-SEEDS>.

	Distinct-1/2	Knowledge Recall	Knowledge Precision	Knowledge F1
Seq2Seq	0.006 / 0.016	0.0047	0.0109	0.0065
Mem-Net	0.010 / 0.030	0.0262	0.0478	0.0327
KG-Net	0.021 / 0.092	0.1605	0.1947	0.1690
RL-DG	0.012 / 0.049	0.0891	0.1079	0.0941
Our Method	0.021 / 0.097	0.2059	0.2167	0.2050

Table 1: Experimental results with automatic measurements, with highest value written in bold.

is increased by 3.6%, which is a significant improvement.

3.2.2 Human Evaluation

Currently, most automatic metrics are not aligned well with human beings in dialogue evaluation (Liu et al., 2016), such as BLEU, ROUGE, etc. In our experiments, extensive evaluations have been carried out with crowd-sourced human beings. With the background knowledge (persona profiles of two participants) and the first start utterance in the testing set, simulated dialogues were generated using each method. There are 8 turns in the simulated conversations (1 start utterance followed by 7 successive generated responses).

Our method is compared with the rest state-of-the-art approaches and each group contains 100 pairs of simulated dialogues, randomly selected from the testing set. For each pair of conversations, they share the same background knowledge and 3 crowd-sourced workers are asked to compare these two simulated conversations at the same time. The human evaluations include the following aspects: (1) Overall refers to the general preference towards the two conversations, with a joint consideration of effective information exchange and coherent interaction. (2) Coverage measures the amount of knowledge expressed during conversations. (3) Concise considers the information repetition and utterance reiteration within conversations. (4) Coherence estimates the consistency and appropriateness within the interaction between participants.

The final comparison results by crowd-sourced workers are determined through majority voting, which are summarized in Table 2. These results demonstrate that our method is consistently and significantly better than the other state-of-the-art approaches.

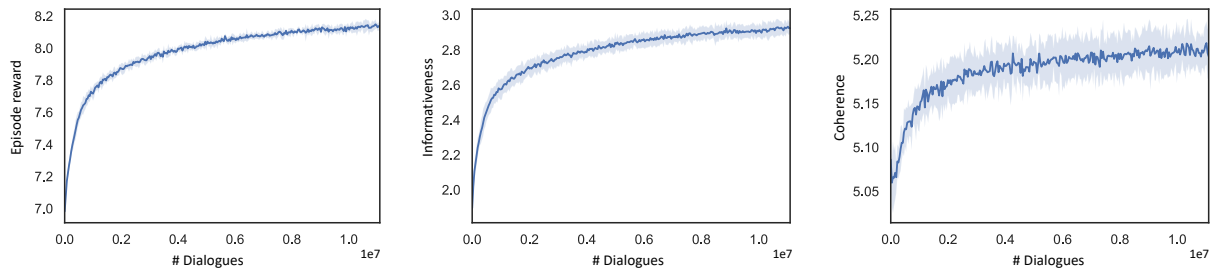


Figure 5: Training curves of reinforcement learning.

	Our Method v.s. Mem-Net			Our Method v.s. KG-Net			Our Method v.s. RL-DG		
	Our Method win	Mem-Net win	Tie	Our Method win	KG-Net win	Tie	Our Method win	RL-DG win	Tie
Overall	83	9	8	65	17	18	79	12	9
Coverage	95	1	4	58	16	26	82	6	12
Concise	78	5	17	47	10	43	77	7	16
Coherence	68	13	19	52	15	33	71	12	17

Table 2: Experimental results with human evaluation, with highest value written in bold.

Personal	Persona 2
i love romance books and movies. i rent an apartment. i want to be a chef. i am scared of dogs. i am a 32 year old single woman.	i am a student. i have a dog named max. my girlfriend lives with me. i enjoy traveling and going on adventures. i am studying to be a lawyer.
Mem-Net	RL-DG
[Start] hello there, how are you? [P1] i am doing well. how are you? [P2] i am good thanks for asking [P1] what do you do for a living? [P2] i am a student. what about you? [P1] i am a student. what do you study? [P2] i am studying dentistry. what about you? [P1] i am studying to be a nurse.	[Start] hello there, how are you? [P1] i am good. just got done walking my dogs. [P2] what do you do for a living? [P1] i am a chef. i am a chef. what do you do for work? [P2] i do not have a job. what do you do for a living? [P1] i am a chef, chef, chef, chef, chef, chef, chef, chef, chef, chef, chef, chef, chef [P2] chef chef chef chef chef? that sounds interesting. do you have a family? [P1] i do not. i have two children. how about you?
KG-Net	Our Method
[Start] hello there, how are you? [P1] i am doing great. just got done with my boyfriend. [P2] i am going to hang out with my friends. [P1] do you have any pets? i have 3 dogs. [P2] i have a dog, his name is bob. [P1] that is a cool name. i love to take my dogs for a walk [P2] do you have a dog? i have a dog. [P1] i do not have a dog. i grew up in a apartment.	[Start] hello there, how are you? [P1] i am ok. just came back from my apartment. [P2] i am studying to be a lawyer. [P1] interesting. i love to read crime novels [P2] i love to go camping sometimes [P1] do you have any pets? i am a stay at home mom. [P2] yes i do. i am a full time student [P1] cool. do you like to read ?

Table 3: Simulated dialogues with the same personas and start utterance.




Figure 6: Visualisation of knowledge utilization in conversations of our method (Upper) and KG-Net (Bottom). Horizontal: background knowledge in the first 12 simulated dialogues, separated by Purple lines. Vertical: knowledge selection probability of each response by one participant.

3.3 Discussions

3.3.1 Case Analysis

Table 3 provides several detailed cases of the simulated dialogues generated by each method, under the same background knowledge (persona profiles) and the start utterance. It can be observed that Mem-Net tends to generate general and fluent responses, like “what about you”, while expresses limited background knowledge. Although informative utterances can be generated by KG-Net, due to a lack of control on information utilization, serious repetition has emerged in the simulated conversation. In addition to redundant responses, another problem with RL-DG is the poor linguistic quality, which might be caused by the decoder update via RL (Lewis et al., 2017; Yarats and Lewis, 2018). Our method is able to generate informative and coherent conversations because the decoder is fixed and only the knowledge selection strategy keeps evolving via compound assessment

Visualization of knowledge utilization in conversations is displayed in Figure 6, where the first 12 simulated dialogues from the testing set are presented. The horizontal axis is the background knowledge in the dialogues, separated by Purple lines. The vertical axis shows the knowledge selection probability $p(z_i|c_t)$ of each utterance, made by one participant in the simulated dialogues (in total 4 utterances). The upper part (our method) demonstrates extensive knowledge coverage, while the bottom part (KG-Net) exhibits repetitive knowledge utilization (highlighted with red circles).

3.3.2 Correlation Analysis

The correlation statistics between automatic metrics (including the distinct-1/2, knowledge-R/P/F1 and our compound reward) and human annotations are provided in Table 4. The Pearson correlation coefficient (Benesty et al., 2009) is estimated using the annotated overall score of our method v.s.

	Distinct-1/2	Knowledge R/P/F1	Compound Reward
Pearson	0.441 / 0.490	0.516 / 0.048 / 0.441	0.623

Table 4: Correlation between automatic metrics and human evaluations, with highest value written in bold.

	Distinct-1/2	Knowledge Recall	Knowledge Precision	Knowledge F1
Lost in Conversation	0.023 / 0.104	0.1010	0.1399	0.1129
Our Method (Transformer)	0.032 / 0.170	0.3098	0.3207	0.3035

Table 5: Comparison with Lost in Conversation, with highest value written in bold.

KG-Net. These results indicate our designed compound reward is aligned better with human beings than commonly used metrics.

3.3.3 Further Evaluation of the Dialogue Strategy

The PersonaChat dataset is also employed by the ConvAI2 challenge (Dinan et al., 2019a), where the team Lost in Conversation obtained the best performance. The network of Lost in Conversation involves 12 transformer layers, which requires extra training data in addition to PersonaChat. For fair comparison, our dialogue strategy is also implemented with the same number of transformer layers and training settings used by Lost in Conversation. The comparison is summarized in Table 5, which verifies the superiority of our proposed method over the advanced transformer network.

4 Related Work

Our work is related with knowledge grounded response generation and multi-turn conversation with reinforcement learning.

As conventional Seq2Seq (Vinyals and Le, 2015) tends to generate general and dull re-

sponses, some knowledge grounded approaches have been introduced to increase the informativeness with extra knowledge. MemNet (Ghazvininejad et al., 2018) encodes factual texts into memory and decodes via attention mechanism for informative generation. CCM (Zhou et al., 2018) relies on structured knowledge to generate rich-information response. In Lian et al. (2019), the posterior distribution is estimated and accurate knowledge is selected to boost informative generation. However, without thorough consideration and control on the knowledge utilization in multi-turn conversations, the above approaches are prone to produce repetitive and incoherent utterances.

The technique of reinforcement learning has been applied to multi-turn dialogue systems in several scenarios. In RL-DG (Li et al., 2016b), three rewards are defined and combined together to boost diverse response generation. Due to a lack of effective control on knowledge utilization, RL-DG is unable to express extensive information during conversations. As RL-DG relies on the reinforcement signal to update all components in the dialogue system, including decoder, it suffers from poor linguistic quality. In Yao et al. (2018), reinforcement learning is employed to plan a cue word (topic) path for a dialogue, where the cue word at t -th turn will assist the corresponding response generation. Different from these chat approaches, our dialogue generation is conducted under the objective of facilitating effective information exchange and letting both participants know more about each. With judiciously design of evaluation metrics, our compound reward is aligned well with human beings and provides meaningful reinforcement signal to evolve the dialogue strategy.

5 Conclusion

In this paper, a novel Generation-Evaluation framework is proposed for informative and coherent multi-turn dialogue generation. Knowledge grounded conversations are generated under the dialogue strategy, which is able to continuously evolve via reinforcement learning with the compound reward. Comprehensive experimental results demonstrate that the proposed method obtains superior performances than the other state-of-the-art methods on both automatic measurements and human evaluations.

In the future, our work can be potentially im-

proved by enriching the assessments with more fine-grained criteria, which can fully integrate turn-level cohesion and dialogue-level coherence. We will also explore to make full use of knowledge to guide the selection of policy strategies for multi-turn conversation.

Acknowledgments

We would like to thank the ACL reviewers for their constructive suggestions and Jinhua Peng, Chaotao Chen, Min Xie for the helpful discussions. This work was supported by the Natural Science Foundation of China (No.61533018).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *International Conference on Learning Representations*.
- Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. In *Noise reduction in speech processing*, pages 1–4. Springer.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2019a. The second conversational intelligence challenge (convai2). *arXiv preprint arXiv:1902.00098*.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019b. Wizard of wikipedia: Knowledge-powered conversational agents. *International Conference on Learning Representations*.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM*

- international conference on Conference on information & knowledge management*, pages 2333–2338. ACM.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016b. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, volume 1, pages 986–995.
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to select knowledge for response generation in dialog systems. *arXiv preprint arXiv:1902.04911*.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1073–1083.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, volume 1, pages 1577–1586.
- Alessandro Sordani, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. 2015. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 553–562. ACM.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 76–85.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 496–505.
- Lili Yao, Ruijian Xu, Chao Li, Dongyan Zhao, and Rui Yan. 2018. Chat more if you like: Dynamic cue words planning to flow longer conversations. *arXiv preprint arXiv:1811.07631*.
- Denis Yarats and Mike Lewis. 2018. Hierarchical text generation and planning for strategic dialogue. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 5591–5599.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 2204–2213.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 654–664.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense knowledge aware conversation generation with graph attention. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4623–4629.