

Improving Response Quality with Backward Reasoning in Open-domain Dialogue Systems

Ziming Li

University of Amsterdam
Amsterdam, The Netherlands
z.li@uva.nl

Julia Kiseleva

Microsoft
Redmond, United States
julia.kiseleva@microsoft.com

Maarten de Rijke

University of Amsterdam
& Ahold Delhaize
Amsterdam, The Netherlands
m.derijke@uva.nl

ABSTRACT

Being able to generate informative and coherent dialogue responses is crucial when designing human-like open-domain dialogue systems. Encoder-decoder-based dialogue models tend to produce generic and dull responses during the decoding step because the most predictable response is likely to be a non-informative response instead of the most suitable one. To alleviate this problem, we propose to train the generation model in a bidirectional manner by adding a backward reasoning step to the vanilla encoder-decoder training. The proposed backward reasoning step pushes the model to produce more informative and coherent content because the forward generation step's output is used to infer the dialogue context in the backward direction. The advantage of our method is that the forward generation and backward reasoning steps are trained simultaneously through the use of a latent variable to facilitate bidirectional optimization. Our method can improve response quality without introducing side information (e.g., a pre-trained topic model). The proposed bidirectional response generation method achieves state-of-the-art performance for response quality.

CCS CONCEPTS

• Information systems → Chat; Question answering.

KEYWORDS

Open-domain dialogue system; response generation

ACM Reference Format:

Ziming Li, Julia Kiseleva, and Maarten de Rijke. 2021. Improving Response Quality with Backward Reasoning in Open-domain Dialogue Systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3404835.3463004>

1 INTRODUCTION

Recently developed end-to-end dialogue systems are trained using large volumes of human-human dialogues to capture underlying interaction patterns [4, 10, 14, 16, 29, 35, 38]. A commonly used

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '21, July 11–15, 2021, Virtual Event, Canada

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8037-9/21/07...\$15.00
<https://doi.org/10.1145/3404835.3463004>

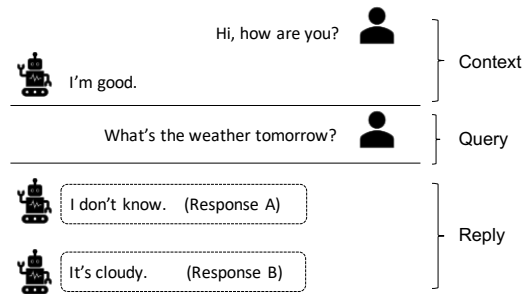


Figure 1: A more informative response (Response B in the figure) can provide information that helps to infer the query content given the dialogue context.

approach to designing data-driven dialogue systems is to use an encoder-decoder framework: feed the dialogue context to the encoder, and let the decoder output an appropriate response. Building on this foundation, different directions have been explored to design dialogue systems that tend to interact with humans in a coherent and engaging manner [2, 15, 19, 31, 34, 37]. However, despite significant advances, there is still room for improvement in the quality of machine-generated responses.

An important problem with encoder-decoder dialogue models is their tendency to generate generic and dull responses, such as "I don't know" or "I'm not sure" [2, 9, 14, 15]. There are two types of methods for dealing with this problem. The first introduces updating signals during training, such as modeling future rewards (e.g., ease of answering) by applying reinforcement learning [15, 19], or bringing variants or adding constraints to the decoding step [2, 14, 31]. The second type holds that, by itself, the dialogue history is not enough for generating high-quality responses, and side information should be taken into account, such as topic information [34, 35] or personal user profiles [37]. Solutions relying on large pre-trained language models, such as DialoGPT [38], can be classified into the second family as well.

In this paper, we propose to train dialogue generation models *bidirectionally* by adding a backward reasoning step to the vanilla encoder-decoder training process. We assume that the information flow in a conversation should be coherent and topic-relevant. Given the dialogue history, neighboring turns are supposed to have a tight topical connection to infer the partial content of one turn given the previous turn *and vice versa*. Inferring the next turn given the (previous) conversation history and the current turn is the traditional take on the dialogue generation task. We extend it by adding one more step: given the dialogue history and the next turn, we aim to infer the content of the current turn. We call the latter step *backward reasoning*. We hypothesize that this can push

the generated response to be more informative and coherent: it is unlikely to infer the dialogue topic given a generic and dull response in the backward direction. An example is shown in Figure 1. Given the dialogue context and *query*,¹ we can predict the reply following a traditional encoder-decoder dialogue generation setup. In contrast, we can infer the content of *query* given the context and reply as long as the reply is informative. Inspired by Zheng et al. [39], we introduce a latent space as a bridge to simultaneously train the encoder-decoder model from two directions. Our experiments demonstrate that the resulting dialogue generation model, called *Mirror*, benefits from this bidirectional training process.

Overall, our work provides the following contributions:

- C1 We introduce a dialogue generation model, *Mirror*, for generating high quality responses in open-domain dialogue systems;
- C2 We define a new way to train dialogue generation models bidirectionally by introducing a latent variable; and
- C3 We obtain improvements in terms of dialogue generation performance with respect to human evaluation on two datasets.

2 RELATED WORK

Conversational scenarios being considered today are increasingly complex, going beyond the ability of rule-based dialogue systems [30]. Ritter et al. [22] propose a data-driven approach to generate responses, building on phrase-based statistical machine translation. Neural network-based models have been studied to generate more informative and interesting responses [23, 26, 29]. Serban et al. [24] introduce latent stochastic variables that span a variable number of time steps to facilitate the generation of long outputs. Deep reinforcement learning methods have also been applied to generate coherent and interesting responses by modeling the future influence of generated responses [15, 19]. Retrieval-based methods are also popular in building dialogue systems by learning a matching model between the context and pre-defined response candidates for response selection [7, 20, 28, 33]. Our work focuses on response *generation* rather than *selection*.

Since encoder-decoder models tend to generate generic and dull responses, Li et al. [14] propose using maximum mutual information as the objective function in neural models to generate more diverse responses. Xing et al. [35] consider incorporating topic information into the encoder-decoder framework to generate informative and interesting responses. To address the dull-response problem, Baheti et al. [2] propose incorporating side information in the form of distributional constraints over the generated responses. Su et al. [27] propose a new perspective to diversify dialogue generation by leveraging non-conversational text. Recently, pre-trained language models, such as GPT-2 [21], Bert [5], XL-Net [36], have been proved effective for a wide range of natural language processing tasks. Several authors make use of pre-trained transformers to attain performance close to humans both in terms of automatic and human evaluation [6, 32, 38]. Though pre-trained language models can perform well for general dialogue generation, they may become less effective without enough data or resources to support these

¹We use *query* to distinguish the current dialogue turn from the context and the response; *query* is not necessarily a real query or question as considered in search or question-answering tasks.

models’ pre-training. In this work, we show the value of developing dialogue generation models with limited data and resources.

The key distinction compared to previous efforts [2, 14] is our work is the first to use the original training dataset through a differentiable backward reasoning step, without external information.

3 METHOD: *MIRROR*

3.1 Problem setting

In many conversational scenarios, the dialogue context is relatively long and contains a lot of information, while the reply (*Response*) is short (and from a different speaker). This makes it difficult to predict the information in the context by only relying on the response in the backward direction. Therefore, we decompose the dialogue context into two different segments: the context c and query x (Figure 1). Assuming that we are predicting the response at turn t in a dialogue, the context c will consist of the dialogue turns from $t - m$ to $t - 2$ and the query x corresponds to turn $t - 1$. Here, we use the term *query* to distinguish the dialogue turn at time step $t - 1$ from the context c and response y ; as explained before, the term *query* should not be confused with a query or question as in search or question-answering tasks. The value m indicates how many dialogue turns we keep in the context c . We use c_{all} to represent the concatenation of c and x , which is also the original context before being decomposed. Our final goal is to predict the response y given dialogue context c and query x .

3.2 Mirror-generative dialogue generation

Shen et al. [25] propose to maximize the conditional log likelihood of generating response y given context c_{all} , $\log p(y | c_{all})$, and they introduce a latent variable z to group different valid responses according to the context c_{all} . The lower bound of $\log p(y | c_{all})$ is given as:

$$\log p(y | c_{all}) \geq \mathbb{E}_{z \sim q_\phi(z | c_{all}, y)} \log p_\theta(y | c_{all}, z) - D_{KL}(q_\phi(z | c_{all}, y) \| p_\theta(z | c_{all})). \quad (1)$$

In Eq. 1, $q_\phi(z | c_{all}, y)$ is the posterior network while $p_\theta(z | c_{all})$ is the prior one.

Instead of maximizing the conditional log likelihood $\log p(y | c_{all})$, we propose to maximize $\log p(x, y | c)$, representing the conditional likelihood that $\langle x, y \rangle$ appears together given dialogue context c . The main assumption underlying this change is that in a conversation, the information flow between neighboring turns should be coherent and relevant, and this connection should be bidirectional. For example, it is not possible to infer what the query is about when a generic and non-informative reply “I don’t know” is given as shown in Figure 1. By taking into account the information flow from two different directions, we hypothesize that we can build a closer connection between the response and the dialogue history and generate more coherent and informative responses. Therefore, we propose to optimize $\log p(x, y | c)$ instead of $\log p(y | c_{all})$.

Following [12, 25], we choose to maximize the variational lower bound of $\log p(x, y | c)$, which is given as:

$$\log p(x, y | c) \geq \mathbb{E}_{z \sim q_\phi(z | c, x, y)} \log p_\theta(x, y | c, z) - D_{KL}(q_\phi(z | c, x, y) \| p_\theta(z | c)), \quad (2)$$




Figure 2: The main architecture of our model, *Mirror*. It consists of three steps: information encoding, latent variable generation, and target decoding.

where z is a shared latent variable between context c , query x and response y . Next, we explain how we optimize a dialogue system by maximizing the lower bound shown in Eq. 2 from two directions.

3.2.1 Forward generation in dialogue generation. With respect to the forward dialogue generation, we interpret the conditional likelihood $\log p_\theta(x, y | c, z)$ in the forward direction:

$$\log p_\theta(x, y | c, z) = \log p_\theta(y | c, z, x) + \log p_\theta(x | c, z). \quad (3)$$

Therefore, we can rewrite Eq. 2 in the forward direction as:

$$\begin{aligned} \log p(x, y | c) &\geq \mathbb{E}_{z \sim q_\phi(z | c, x, y)} [\log p_\theta(y | c, z, x) + \log p_\theta(x | c, z)] \\ &\quad - D_{KL}(q_\phi(z | c, x, y) \| p_\theta(z | c)). \end{aligned} \quad (4)$$

We introduce $q_\phi(z | c, x, y)$ as the posterior network, also referred to as the recognition net, and $p_\theta(z | c)$ as the prior network.

3.2.2 Backward reasoning in dialogue generation. As in the forward direction, if we decompose the conditional likelihood $\log p_\theta(x, y | c, z)$ in the backward direction, we can rewrite Eq. 2 as:

$$\begin{aligned} \log p(x, y | c) &\geq \mathbb{E}_{z \sim q_\phi(z | c, x, y)} [\log p_\theta(x | c, y, z) + \log p_\theta(y | c, z)] \\ &\quad - D_{KL}(q_\phi(z | c, x, y) \| p_\theta(z | c)). \end{aligned} \quad (5)$$

3.2.3 Optimizing dialogue systems bidirectionally. Since the variable z is sampled from the shared latent space between forward generation and backward reasoning steps, we can regard z as a bridge to connect the training in two different direction and this opens the possibility to train dialogue models effectively. By merging Eq. 4 and Eq 5, we can rewrite the lower bound Eq. 2 as:

$$\begin{aligned} \log p(x, y | c) &\geq \mathbb{E}_{z \sim q_\phi(z | c, x, y)} \left[\frac{1}{2} \log p_\theta(x | c, z, y) \right. \\ &\quad + \frac{1}{2} \log p_\theta(y | c, z) + \frac{1}{2} \log p_\theta(y | c, z, x) \\ &\quad \left. + \frac{1}{2} \log p_\theta(x | c, z) - D_{KL}(q_\phi(z | c, x, y) \| p_\theta(z | c)) \right] \\ &= L(c, x, y; \theta, \phi), \end{aligned} \quad (6)$$

which is the final loss function for our dialogue generation model.

3.2.4 Model architecture. The complete architecture of the proposed joint training process is shown in Figure 2. It consists of three steps: (1) information encoding, (2) latent variable generation, and (3) target decoding. With respect to the information encoding step, we utilize a context encoder Enc_{ctx} to compress the dialogue

context c while an utterance encoder Enc_{utt} is used to compress the query x and response y , respectively. To model the latent variable z , we assume z follows the multivariate normal distribution, the posterior network $q_\phi(z | c, x, y) \sim N(\mu, \sigma^2 I)$ and the prior network $p_\theta(z | c) \sim N(\mu', \sigma'^2 I)$. Then, by applying the reparameterization trick [12], we can sample a latent variable z from the estimated posterior distribution $N(\mu, \sigma^2 I)$. During testing, we use the prior distribution $N(\mu', \sigma'^2 I)$ to generate the variable z . The KL-divergence distance is applied to encourage the approximated posterior $N(\mu, \sigma^2 I)$ to be close to the prior $N(\mu', \sigma'^2 I)$. According to Eq. 6, the decoding step in the right side of Figure 2 consists of four independent decoders, Dec_1 , Dec_2 , Dec_3 , and Dec_4 , corresponding to $\log p(y | c, z, x)$, $\log p(x | c, z)$, $\log p(x | c, z, y)$ and $\log p(y | c, z)$, respectively. Decoder Dec_1 is used to generate the final response during the testing stage. To make full use of the variable z , we attach it to the input of each decoding step. Since we have the shared latent vector z as a bridge, training for the two directions is not independent, and updating one direction will definitely improve the other direction as well. In the end, both directions will contribute to the final dialogue generation process.

4 EXPERIMENTAL SETUP

4.1 Datasets

We use two datasets. First, the MovieTriples dataset [23] has been developed by expanding and preprocessing the Movie-Dic corpus [3] of film transcripts and each dialogue consists of 3 turns between two speakers. We regard the first turn as the dialogue context while the second and third one as the query and response, respectively. In the final dataset, there are around 166k dialogues in the training set, 21k in the validation set and 20k in the test set. In terms of the vocabulary table size, we set it to the top 20k most frequent words in the dataset.

Second, the DailyDialog dataset [18] is a high-quality multi-turn dialogue dataset. We split the dialogues in the original dataset into shorter dialogues by every three turns as a new dialogue. The last turn is used as the target response and the first as the context and the third one as the query. After preprocessing, we have 65k, 6k, and 6k dialogs in the training, testing and validation sets, respectively. We limit the vocabulary table size to the top 20k most frequent words for the DailyDialog dataset.

4.2 Baselines

- Seq2SeqAtt** This is a LSTM-based [8] dialogue generation model with attention mechanism [1].
- HRED** This method [23] uses a hierarchical recurrent encoder-decoder to sequentially generate the tokens in the replies.
- VHRED** This extension of HRED incorporates a stochastic latent variable to explicitly model generative processes that possess multiple levels of variability [24]. This is also the model trained with Eq. 1.
- MMI** This method first generates response candidates on a Seq2Seq model trained in the direction of context-to-target, $P(y | c, x)$, then re-ranks them using a separately trained Seq2Seq model in the direction of target-to-context, $P(x | y)$, to maximize the mutual information [14].

DC This method incorporates side information in the form of distributional constraints, including topic constraints and semantic constraints [2].

DC-MMI This method is a combination of **MMI** and **DC**, where the decoding step takes into account mutual information together with the proposed distribution constraints in the method **DC**.

4.3 Training details

We implement our model, *Mirror*², with PyTorch in the OpenNMT framework [13]. The utterance encoder is a two-layer LSTM [8] and the dimension is 1,000. The context encoder has the same architecture as the utterance encoder but the parameters are not shared. The four decoders have the same design but independent parameters, and each one is a two-layer LSTM with 1,000 dimensions. In terms of the dimension of the hidden vector z , we set it to 160 for the DailyDialog dataset while 100 for MovieTriples. The word embedding size is 200 for both datasets. We use Adam [11] as the optimizer. The initial learning rate is 0.001 and learning rate decay is applied to stabilize the training process.

4.4 Evaluation

We conduct a human evaluation on Amazon MTurk guided by [17]. For each two-way comparison of dialogue responses (against *Mirror*), we ask annotators to judge which of two responses is more appropriate given the context. For each method pair (*Mirror*, *Baseline*) and each dataset, we randomly sample 200 dialogues from the test datasets; each pair of responses is annotated by 3 annotators.

5 RESULTS AND ANALYSIS

In Table 1, we show performance comparisons between *Mirror* and other baselines on two different datasets. According to Table 1(top), it is somewhat unexpected to see that HRED can achieve such close performance compared to *Mirror* on DailyDialog, given its main architecture is a hierarchical encoder-decoder model. We randomly sample some dialogue pairs for which HRED outperforms *Mirror* to see why annotators prefer HRED over *Mirror*. For many of these cases, *Mirror* fails to generate appropriate responses, while HRED returns generic but still acceptable responses given the context. When we have the back reasoning step in *Mirror*, we expect that it will lead to more informative generations. Still, it also increases the risk of generating responses with incorrect syntax or relevant but inappropriate responses. A possible reason for the latter is that the backward reasoning step has dominated the joint training process, which can degenerate the forward generation performance.

The performance gap between *Mirror* and all approaches (including HRED) is large on the DailyDialog dataset (see Table 1(bottom)). Due to space limitations, we only present one dialogue example in Table 2. The example is a typical case of why the response generated by DC has high embedding scores, but the human evaluation result is not promising. In this example, the response from DC has high semantic similarity with the context because of words like “ask you”, “apartment”, and “questions”. However, it cannot be regarded as an appropriate and meaningful response in the given context. Comparing *Mirror* with methods that have use **MMI** (**MMI**, **DC-MMI**), the performance gap is relatively small. This is evidence showing

²Codebase: <https://github.com/cszmli/mirror-sigir>

Table 1: Human evaluation using the MovieTriple and Daily-Dialog datasets.

	Method pair	Wins	Losses	Ties
(a) <i>MovieTriple</i>	Mirror vs. Seq2SeqAttn	0.53	0.37	0.10
	Mirror vs. HRED	0.41	0.40	0.19
	Mirror vs. VHRED	0.45	0.38	0.17
	Mirror vs. MMI	0.48	0.42	0.10
	Mirror vs. DC	0.50	0.33	0.17
	Mirror vs. DC-MMI	0.39	0.35	0.26
(b) <i>DailyDialog</i>	Mirror vs. Seq2SeqAttn	0.50	0.26	0.24
	Mirror vs. HRED	0.49	0.32	0.19
	Mirror vs. VHRED	0.48	0.37	0.15
	Mirror vs. MMI	0.40	0.34	0.26
	Mirror vs. DC	0.45	0.38	0.17
	Mirror vs. DC-MMI	0.47	0.35	0.18

the effectiveness of maximizing mutual information in improving the response quality. The *Mirror* method can be treated as a way to maximize mutual information implicitly. The advantage is that we can train dialogue models in two directions simultaneously.

Table 2: Example generated responses by different models when the dialogue context is given.

Context
<p>Speaker A: here 's my license .</p> <p>Speaker B: i 'm afraid i 'm going to have to ask you to remain in the apartment . the narcotics squad will be arriving any moment now . they want to ask you a few questions .</p>
Response
<p>Reference: squad ? what do they want with me ? i don 't even use aspirin !</p> <p>Seq2Seq: no .</p> <p>HRED: i don 't think so .</p> <p>VHRED: oh , i 'm sorry .</p> <p>MMI: i 'm sorry . i 'm sorry . i don 't know what you 're talking about . i don 't know what i 'm afraid of .</p> <p>DC: i 'm not going to ask you . but he will be in the apartment for a moment - and we can have some questions with that one of them !</p> <p>DC-MMI: i 'm going to ask you .</p> <p>Mirror: well , i 'm sure they 'll have to wait .</p>

6 CONCLUSION AND FUTURE WORK

We have presented a novel approach to generating informative and coherent responses in open-domain dialogue systems, called *Mirror*. First, we reformulate the original response generation task from two sides: context and response, to three sides: context, query, and response. Given the dialogue context and query, predicting the response is exactly like the traditional dialogue generation setup. Thus, *Mirror* has one more step: inferring the query given the dialogue context and response. By incorporating the backward reasoning step, we implicitly push the model to generate responses that have closer connections with the dialogue history. By conducting experiments on two datasets, we have demonstrated that *Mirror* improves the response quality compared to several competitive baselines without incorporating additional sources of information, which comes with additional computational costs and complexity. For future work, *Mirror*'s bidirectional training approach can be generalized to other domains, such as task-oriented dialogue systems and question-answering tasks.

REFERENCES

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv preprint arXiv:1409.0473* (2014).
- [2] Ashutosh Baheti, Alan Ritter, Jiwei Li, and Bill Dolan. 2018. Generating More Interesting Responses in Neural Conversation Models with Distributional Constraints. *arXiv preprint arXiv:1809.01215* (2018).
- [3] Rafael E Banchs. 2012. Movie-DiC: A Movie Dialogue Corpus for Research and Development. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*. Association for Computational Linguistics, 203–207.
- [4] Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2019. Plato: Pre-trained Dialogue Generation Model with Discrete Latent Variable. *arXiv preprint arXiv:1910.07931* (2019).
- [5] 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* (2018).
- [6] Sergey Golovanov, Rauf Kurbanov, Sergey Nikolenko, Kyryl Truskovskiy, Alexander Tselousov, and Thomas Wolf. 2019. Large-scale Transfer Learning for Natural Language Generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 6053–6058.
- [7] Jia-Chen Gu, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, and Xiaodan Zhu. 2020. Speaker-aware Bert for Multi-turn Response Selection in Retrieval-based Chatbots. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2041–2044.
- [8] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-term Memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [9] Shaojie Jiang and Maarten de Rijke. 2018. Why are Sequence-to-Sequence Models So Dull? Understanding the Low-Diversity Problem of Chatbots. In *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*. ACL.
- [10] Chandra Khatri, Behnam Hedayatnia, Anu Venkatesh, Jeff Nunn, Yi Pan, Qing Liu, Han Song, Anna Gottardi, Sanjeev Kwatra, Sanju Panchoi, et al. 2018. Advancing the State of the Art in Open Domain Dialog Systems through the Alexa Prize. *arXiv preprint arXiv:1812.10757* (2018).
- [11] Diederik P Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [12] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [13] Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. OpenNMT: Open-Source Toolkit for Neural Machine Translation. In *Proc. ACL*.
- [14] Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A Diversity-promoting Objective Function for Neural Conversation Models. *arXiv preprint arXiv:1510.03055* (2015).
- [15] Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep Reinforcement Learning for Dialogue Generation. *arXiv preprint arXiv:1606.01541* (2016).
- [16] Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial Learning for Neural Dialogue Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 2157–2169.
- [17] Margaret Li, Jason Weston, and Stephen Roller. 2019. Acute-eval: Improved Dialogue Evaluation with Optimized Questions and Multi-turn comparisons. *arXiv preprint arXiv:1909.03087* (2019).
- [18] Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A Manually Labelled Multi-turn Dialogue Dataset. *arXiv preprint arXiv:1710.03957* (2017).
- [19] Ziming Li, Julia Kiseleva, and Maarten de Rijke. 2019. Dialogue Generation: From Imitation Learning to Inverse Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6722–6729.
- [20] Lisong Qiu, Yingwai Shiu, Pingping Lin, Ruihua Song, Yue Liu, Dongyan Zhao, and Rui Yan. 2020. What If Bots Feel Moods?. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1161–1170.
- [21] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. *OpenAI blog* 1, 8 (2019), 9.
- [22] Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven Response Generation in Social Media. In *Proceedings of the conference on empirical methods in natural language processing*. Association for Computational Linguistics, 583–593.
- [23] Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C Courville, and Joelle Pineau. 2016. Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models. In *AAAI*, Vol. 16. 3776–3784.
- [24] Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A Hierarchical Latent Variable Encoder-decoder Model for Generating Dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [25] Xiaoyu Shen, Hui Su, Yanran Li, Wenjie Li, Shuzi Niu, Yang Zhao, Akiko Aizawa, and Guoping Long. 2017. A Conditional Variational Framework for Dialog Generation. *arXiv preprint arXiv:1705.00316* (2017).
- [26] Alessandro Sordani, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. A Neural Network Approach to Context-sensitive Generation of Conversational Responses. *arXiv preprint arXiv:1506.06714* (2015).
- [27] Hui Su, Xiaoyu Shen, Sanqiang Zhao, Xiao Zhou, Pengwei Hu, Randy Zhong, Cheng Niu, and Jie Zhou. 2020. Diversifying Dialogue Generation with Non-conversational Text. *arXiv preprint arXiv:2005.04346* (2020).
- [28] Chongyang Tao, Wei Wu, Can Xu, Wenpeng Hu, Dongyan Zhao, and Rui Yan. 2019. Multi-representation Fusion Network for Multi-turn Response Selection in Retrieval-based Chatbots. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 267–275.
- [29] Oriol Vinyals and Quoc Le. 2015. A Neural Conversational Model. *arXiv preprint arXiv:1506.05869* (2015).
- [30] Joseph Weizenbaum. 1966. ELIZA—A Computer Program for the Study of Natural Language Communication between Man and Machine. *Commun. ACM* 9, 1 (1966), 36–45.
- [31] Sam Wiseman and Alexander M Rush. 2016. Sequence-to-sequence Learning as Beam-search Optimization. *arXiv preprint arXiv:1606.02960* (2016).
- [32] Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A Transfer Learning Approach for Neural Network based Conversational Agents. *arXiv preprint arXiv:1901.08149* (2019).
- [33] Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2016. Sequential Matching Network: A New Architecture for Multi-turn Response Selection in Retrieval-based Chatbots. *arXiv preprint arXiv:1612.01627* (2016).
- [34] Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2016. Topic Augmented Neural Response Generation with a Joint Attention Mechanism. *arXiv preprint arXiv:1606.08340* 2, 2 (2016).
- [35] Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic Aware Neural Response Generation. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [36] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized Autoregressive Pretraining for Language Understanding. *arXiv preprint arXiv:1906.08237* (2019).
- [37] 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? *arXiv preprint arXiv:1801.07243* (2018).
- [38] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2019. Dialogpt: Large-scale Generative Pre-training for Conversational Response Generation. *arXiv preprint arXiv:1911.00536* (2019).
- [39] Zaixiang Zheng, Hao Zhou, Shujian Huang, Lei Li, Xin-Yu Dai, and Jiajun Chen. 2019. Mirror-Generative Neural Machine Translation. In *International Conference on Learning Representations*.