

GTM: A Generative Triple-Wise Model for Conversational Question Generation

Lei Shen^{1,2} Fandong Meng³ Jinchao Zhang³ Yang Feng^{1,2*} Jie Zhou³

¹Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

³Pattern Recognition Center, WeChat AI, Tencent Inc, China

shenleil7z@ict.ac.cn, {fandongmeng, dayerzhang}@tencent.com

fengyang@ict.ac.cn, withtomzhou@tencent.com

Abstract

Generating some appealing questions in open-domain conversations is an effective way to improve human-machine interactions and lead the topic to a broader or deeper direction. To avoid dull or deviated questions, some researchers tried to utilize answer, the “future” information, to guide question generation. However, they separate a post-question-answer (PQA) triple into two parts: post-question (PQ) and question-answer (QA) pairs, which may hurt the overall coherence. Besides, the QA relationship is modeled as a one-to-one mapping that is not reasonable in open-domain conversations. To tackle these problems, we propose a generative triple-wise model with hierarchical variations for open-domain conversational question generation (CQG). Latent variables in three hierarchies are used to represent the shared background of a triple and one-to-many semantic mappings in both PQ and QA pairs. Experimental results on a large-scale CQG dataset show that our method significantly improves the quality of questions in terms of fluency, coherence and diversity over competitive baselines.

1 Introduction

Questioning in open-domain dialogue systems is indispensable since a good system should have the ability to well interact with users by not only responding but also asking (Li et al., 2017). Besides, raising questions is a proactive way to guide users to go deeper and further into conversations (Yu et al., 2016). Therefore, the ultimate goal of open-domain conversational question generation (CQG) is to enhance the interactivity and maintain the continuity of a conversation (Wang et al., 2018).

Joint work with Pattern Recognition Center, WeChat AI, Tencent Inc, China. *Yang Feng is the corresponding author.

Post:

I ate out with my friends this evening.

Question Candidates:

Q1.1: Which restaurant did you go?

Q1.2: Where did you eat?

Q2.1: What food did you eat?

Q2.2: Did you eat something special?

Q3: What do you mean?

Q4: How about drinking together?

Answer Candidates:

A1: We went to an Insta-famous cafeteria.

A2: We ate steak and pasta.

Table 1: An example of CQG task which is talking about a person’s eating activity. There are one-to-many mappings in both PQ and QA pairs. The content of each meaningful and relevant question (Q1.1 to Q2.2) is decided by its post and answer. Q3 (dull) and Q4 (deviated) are generated given only the post.

CQG differs fundamentally from traditional question generation (TQG) (Zhou et al., 2019; Kim et al., 2019; Li et al., 2019) that generates a question given a sentence/paragraph/passage and a specified answer within it. While in CQG, an answer always follows the to-be-generated question, and is unavailable during inference (Wang et al., 2019). At the same time, each utterance in open-domain scenario is casual and can be followed by several appropriate sentences, i.e., one-to-many mapping (Gao et al., 2019; Chen et al., 2019).

At first, the input information of CQG was mainly a given post (Wang et al., 2018; Hu et al., 2018), and the generated questions were usually dull or deviated (Q3 and Q4 in Table 1). Based on the observation that an answer has strong relevance to its question and post, Wang et al. (2019) tried to integrate answer into the question generation process. They applied a reinforcement learning framework that firstly generated a question given the post, and then used a pre-trained matching model to estimate the relevance score (reward) between

answer and generated question. This method separates a post-question-answer (PQA) triple into post-question (PQ) and question-answer (QA) pairs rather than considering the triple as a whole and modeling the overall coherence. Furthermore, the training process of the matching model only utilizes one-to-one relation of each QA pair and neglects the one-to-many mapping feature.

An open-domain PQA often takes place under a background that can be inferred from all utterances in the triple and help enhance the overall coherence. When it comes to the semantic relationship in each triple, the content of a specific question is under the control of its post and answer (Lee et al., 2020). Meanwhile, either a post or an answer could correspond to several meaningful questions. As shown in Table 1, the triple is about a person’s eating activity (the background of the entire conversation). There are one-to-many mappings in both PQ and QA pairs that construct different meaningful combinations, such as P-Q1.1-A1, P-Q1.2-A1, P-Q2.1-A2 and P-Q2.2-A2. An answer connects tightly to both its post and question, and in turn helps decide the expression of a question.

On these grounds, we propose a generative triple-wise model (GTM) for CQG. Specifically, we firstly introduce a triple-level variable to capture the shared background among PQA. Then, two separate variables conditioned on the triple-level variable are used to represent the latent space for question and answer, and the question variable is also dependent on the answer one. During training, the latent variables are constrained to reconstruct both the original question and answer according to the hierarchical structure we define, making sure the triple-wise relationship flows through the latent variables without any loss. For the question generation process, we sample the triple-level and answer variable given a post, then obtain the question variable conditioned on them, and finally generate a question based on the post, triple-level and question variables. Experimental results on a large-scale CQG dataset show that GTM can generate more fluent, coherent, and intriguing questions for open-domain conversations.

The main contribution is threefold:

- To generate coherent and informative questions in the CQG task, we propose a generative triple-wise model that models the semantic relationship of a triple in three levels: PQA, PQ, and QA.




Figure 1: The graphical representation of GTM for training process. z^t is used to capture the shared background among PQA, while z^q and z^a are used to model the diversity in PQ and QA pairs. Solid arrows illustrate the generation of q , a (not used in inference), and qt , while dashed arrows are for posterior distributions of latent variables.

- Our variational hierarchical structure can not only utilize the “future” information (answer), but also capture one-to-many mappings in PQ and QA, which matches the open-domain scenario well.
- Experimental results on a large-scale CQG corpus show that our method significantly outperforms the state-of-the-art baselines in both automatic and human evaluations.

2 Proposed Model

Given a post as the input, the goal of CQG is to generate the corresponding question. Following the work of Zhao et al. (2017) and Wang et al. (2019), we leverage the question type qt to control the generated question, and take advantage of the answer information a to improve coherence. In training set, each conversation is represented as $\{p, q, qt, a\}$, consisting of post $p = \{p_i\}_{i=1}^{|p|}$, question $q = \{q_i\}_{i=1}^{|q|}$ with its question type qt , and answer $a = \{a_i\}_{i=1}^{|a|}$.

2.1 Overview

The graphical model of GTM for training process is shown in Figure 1. θ , φ , and ϕ are used to denote parameters of generation, prior, and recognition network, respectively. We integrate answer generation to assist question generation with hierarchical latent variables. Firstly, a triple-level variable z^t is imported to capture the shared background and

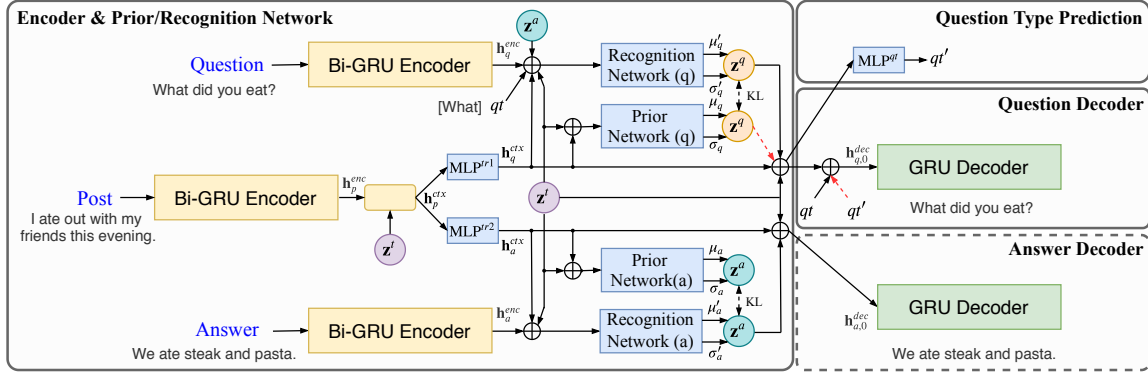


Figure 2: The architecture of GTM. \oplus denotes the concatenation operation. In training process, latent variables obtained from recognition networks and the real question type qt are used for decoding. Red dashed arrows refer to inference process, in which we get latent variables from prior networks, and the predicted question type qt' is fed into the question decoder. The answer decoder is only utilized during training to assist the triple-wise modeling.

is inferred from PQA utterances. Then answer latent variable \mathbf{z}^a and question latent variable \mathbf{z}^q are sampled from Gaussian distributions conditioned on both post and \mathbf{z}^t . To ensure that the question is controlled by answer, \mathbf{z}^q is also dependent on \mathbf{z}^a .

2.2 Input Representation

We use a bidirectional GRU (Cho et al., 2014) as encoder to capture the semantic representation of each utterance. Take post \mathbf{p} as an example. Each word in \mathbf{p} is firstly encoded into its embedding vector. The GRU then computes forward hidden states $\{\vec{h}_i\}_{i=1}^{|\mathbf{p}|}$ and backward hidden states $\{\overleftarrow{h}_i\}_{i=1}^{|\mathbf{p}|}$:

$$\begin{aligned}\vec{h}_i &= \overrightarrow{\text{GRU}}(e_{p_i}, \vec{h}_{i-1}), \\ \overleftarrow{h}_i &= \overleftarrow{\text{GRU}}(e_{p_i}, \overleftarrow{h}_{i+1}),\end{aligned}$$

where e_{p_i} is employed to represent the embedding vector of word p_i . We finally get the post representation by concatenating the last hidden states of two directions $\mathbf{h}_p^{\text{enc}} = [\vec{h}_{|\mathbf{p}|}; \overleftarrow{h}_1]$. Similarly, we can obtain representations of question \mathbf{q} and answer \mathbf{a} , denoted as $\mathbf{h}_q^{\text{enc}}$ and $\mathbf{h}_a^{\text{enc}}$, respectively.

The question type qt is represented by a real-valued, low dimensional vector \mathbf{v}_{qt} which is updated during training and is regarded as a linguistic feature that benefits the training of latent variables (Zhao et al., 2017). We use the actual question type qt during training to provide the information of interrogative words that is the most important feature to distinguish question types.

2.3 Triple-level Latent Variable

To capture the shared background of entire triple, we introduce a triple-level latent variable \mathbf{z}^t that

is inferred from PQA utterances and is in turn responsible for generating the whole triple. Inspired by Park et al. (2018), we use a standard Gaussian distribution as the prior distribution of \mathbf{z}^t :

$$p_\varphi(\mathbf{z}^t) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I}),$$

where \mathbf{I} represents the identity matrix.

For the inference of \mathbf{z}^t in training set, we consider three utterance representations $\mathbf{h}_p^{\text{enc}}$, $\mathbf{h}_q^{\text{enc}}$ and $\mathbf{h}_a^{\text{enc}}$ as a sequence, and use a bidirectional GRU to take individual representation as the input of each time step. The triple representation \mathbf{h}^t is obtained by concatenating the last hidden states of both directions. Then, \mathbf{z}^t is sampled from:

$$\begin{aligned}q_\phi(\mathbf{z}^t|\mathbf{p}, \mathbf{q}, \mathbf{a}) &= \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}^t, \boldsymbol{\sigma}^t\mathbf{I}), \\ \boldsymbol{\mu}^t &= \text{MLP}_\phi^t(\mathbf{h}^t), \\ \boldsymbol{\sigma}^t &= \text{softplus}(\text{MLP}_\phi^t(\mathbf{h}^t)),\end{aligned}$$

where $\text{MLP}(\cdot)$ is a feed-forward network, and softplus function is a smooth approximation to ReLU and can be used to ensure positiveness (Park et al., 2018; Serban et al., 2017).

2.4 One-to-many Mappings

After obtaining \mathbf{z}^t , we use a GRU f to get a vector $\mathbf{h}_p^{\text{ctx}}$ for connecting \mathbf{p} and \mathbf{q}/\mathbf{a} . $\mathbf{h}_p^{\text{ctx}}$ is then transformed to $\mathbf{h}_q^{\text{ctx}}$ and $\mathbf{h}_a^{\text{ctx}}$ that are used in prior and recognition networks for \mathbf{z}^q and \mathbf{z}^a :

$$\begin{aligned}\mathbf{h}_p^{\text{ctx}} &= f(\mathbf{z}^t, \mathbf{h}_p^{\text{enc}}), \\ \mathbf{h}_q^{\text{ctx}} &= \text{MLP}_\theta^{\text{tr}1}(\mathbf{h}_p^{\text{ctx}}), \\ \mathbf{h}_a^{\text{ctx}} &= \text{MLP}_\theta^{\text{tr}2}(\mathbf{h}_p^{\text{ctx}}).\end{aligned}$$

To model one-to-many mappings in PQ and QA pairs under the control of \mathbf{z}^t , we design two utterance-level variables, \mathbf{z}^q and \mathbf{z}^a , to represent latent spaces of question and answer. We define the prior and posterior distributions of \mathbf{z}^a as follows:

$$\begin{aligned} p_\varphi(\mathbf{z}^a|\mathbf{p}, \mathbf{z}^t) &= \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_a, \boldsymbol{\sigma}_a\mathbf{I}), \\ q_\phi(\mathbf{z}^a|\mathbf{p}, \mathbf{z}^t, \mathbf{a}) &= \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}'_a, \boldsymbol{\sigma}'_a\mathbf{I}), \end{aligned}$$

where $\boldsymbol{\mu}_a$, $\boldsymbol{\sigma}_a$, $\boldsymbol{\mu}'_a$, and $\boldsymbol{\sigma}'_a$, the parameters of two Gaussian distributions, are calculated as:

$$\begin{aligned} \boldsymbol{\mu}_a &= \text{MLP}_\varphi^a([\mathbf{h}_a^{ctx}; \mathbf{z}^t]), \\ \boldsymbol{\sigma}_a &= \text{softplus}(\text{MLP}_\varphi^a([\mathbf{h}_a^{ctx}; \mathbf{z}^t])), \end{aligned}$$

$$\begin{aligned} \boldsymbol{\mu}'_a &= \text{MLP}_\phi^a([\mathbf{h}_a^{ctx}; \mathbf{z}^t; \mathbf{h}_a^{enc}], \\ \boldsymbol{\sigma}'_a &= \text{softplus}(\text{MLP}_\phi^a([\mathbf{h}_a^{ctx}; \mathbf{z}^t; \mathbf{h}_a^{enc}])). \end{aligned}$$

To make sure the content of question is also decided by answer and improve their relatedness, we import \mathbf{z}^a into \mathbf{z}^q space. The prior and posterior distributions of \mathbf{z}^q are computed as follows:

$$\begin{aligned} p_\varphi(\mathbf{z}^q|\mathbf{p}, \mathbf{z}^t, \mathbf{z}^a) &= \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_q, \boldsymbol{\sigma}_q\mathbf{I}), \\ q_\phi(\mathbf{z}^q|\mathbf{p}, \mathbf{z}^t, \mathbf{q}, qt, \mathbf{z}^a) &= \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}'_q, \boldsymbol{\sigma}'_q\mathbf{I}), \end{aligned}$$

where $\boldsymbol{\mu}_q$, $\boldsymbol{\sigma}_q$, $\boldsymbol{\mu}'_q$, and $\boldsymbol{\sigma}'_q$ are calculated as:

$$\begin{aligned} \boldsymbol{\mu}_q &= \text{MLP}_\varphi^q([\mathbf{h}_q^{ctx}; \mathbf{z}^t; \mathbf{z}^a]), \\ \boldsymbol{\sigma}_q &= \text{softplus}(\text{MLP}_\varphi^q([\mathbf{h}_q^{ctx}; \mathbf{z}^t; \mathbf{z}^a])), \\ \boldsymbol{\mu}'_q &= \text{MLP}_\phi^q([\mathbf{h}_q^{ctx}; \mathbf{z}^t; \mathbf{h}_q^{enc}; \mathbf{v}_{qt}; \mathbf{z}^a]), \\ \boldsymbol{\sigma}'_q &= \text{softplus}(\text{MLP}_\phi^q([\mathbf{h}_q^{ctx}; \mathbf{z}^t; \mathbf{h}_q^{enc}; \mathbf{v}_{qt}; \mathbf{z}^a])). \end{aligned}$$

2.5 Question Generation Network

Following the work of [Zhao et al. \(2017\)](#) and [Wang et al. \(2019\)](#), a question type prediction network MLP^{qt} is introduced to approximate $p_\theta(qt|\mathbf{z}^q, \mathbf{z}^t, \mathbf{p})$ in training process and produces question type qt' during inference.

As shown in [Figure 2](#), there are two decoders in our model, one is for answer generation that is an auxiliary task and only exists in the training process, and the other is for desired question generation. The question decoder employs a variant of GRU that takes the concatenation result of \mathbf{z}^q , \mathbf{z}^t , \mathbf{h}_q^{ctx} , and qt as initial state \mathbf{s}_0 , i.e., $\mathbf{s}_0 = [\mathbf{z}^q; \mathbf{z}^t; \mathbf{h}_q^{ctx}; qt]$. For each time step j , it calculates the context vector \mathbf{c}_j following [Bahdanau](#)

[et al. \(2015\)](#), and computes the probability distribution $p_\theta(\mathbf{q}|\mathbf{z}^q, \mathbf{z}^t, \mathbf{p}, qt)$ over all words in the vocabulary:

$$\begin{aligned} \mathbf{s}_j &= \text{GRU}(\mathbf{e}_{j-1}, \mathbf{s}_{j-1}, \mathbf{c}_j) \\ \tilde{\mathbf{s}}_j &= \text{MLP}([\mathbf{e}_{j-1}; \mathbf{c}_j; \mathbf{s}_j]), \\ p_\theta(\mathbf{q}_j|\mathbf{q}_{<j}, \mathbf{z}^q, \mathbf{z}^t, \mathbf{p}, qt) &= \text{softmax}(\mathbf{W}_o\tilde{\mathbf{s}}_j), \end{aligned}$$

where \mathbf{e}_{j-1} represents the embedding vector of the $(j-1)$ -th question word. Similarly, the answer decoder receives the concatenation result of \mathbf{z}^a , \mathbf{z}^t , and \mathbf{h}_a^{ctx} as initial state to approximate the probability $p_\theta(\mathbf{a}|\mathbf{z}^a, \mathbf{z}^t, \mathbf{p})$.

2.6 Training and Inference

Importantly, our model GTM is trained to maximize the log-likelihood of the joint probability $p(\mathbf{p}, \mathbf{q}, \mathbf{a}, qt)$:

$$\log p(\mathbf{p}, \mathbf{q}, \mathbf{a}, qt) = \log \int_{\mathbf{z}^t} p(\mathbf{p}, \mathbf{q}, \mathbf{a}, qt, \mathbf{z}^t).$$

However, the optimization function is not directly tractable. Inspired by [Serban et al. \(2017\)](#) and [Park et al. \(2018\)](#), we convert it to the following objective that is based on the evidence lower bound and needs to be maximized in training process:

$$\begin{aligned} \mathcal{L}_{\text{GTM}} &= \\ &- KL(q_\phi(\mathbf{z}^t|\mathbf{p}, \mathbf{q}, \mathbf{a})||p_\varphi(\mathbf{z}^t)) \\ &- KL(q_\phi(\mathbf{z}^a|\mathbf{p}, \mathbf{z}^t, \mathbf{a})||p_\varphi(\mathbf{z}^a|\mathbf{p}, \mathbf{z}^t)) \\ &- KL(q_\phi(\mathbf{z}^q|\mathbf{p}, \mathbf{z}^t, \mathbf{q}, qt, \mathbf{z}^a)||p_\varphi(\mathbf{z}^q|\mathbf{p}, \mathbf{z}^t, \mathbf{z}^a)) \\ &+ \mathbb{E}_{\mathbf{z}^a, \mathbf{z}^t \sim q_\phi} [\log p_\theta(\mathbf{a}|\mathbf{z}^a, \mathbf{z}^t, \mathbf{p})] \\ &+ \mathbb{E}_{\mathbf{z}^q, \mathbf{z}^t \sim q_\phi} [\log p_\theta(\mathbf{q}|\mathbf{z}^q, \mathbf{z}^t, \mathbf{p}, qt)] \\ &+ \mathbb{E}_{\mathbf{z}^q, \mathbf{z}^t \sim q_\phi} [\log p_\theta(qt|\mathbf{z}^q, \mathbf{z}^t, \mathbf{p})]. \end{aligned}$$

The objective consists of two parts: the variational lower bound (the first five lines) and question type prediction accuracy (the last line). Meanwhile, the variational lower bound includes the reconstruction terms and KL divergence terms based on three hierarchical latent variables. The gradients to the prior and recognition networks can be estimated using the reparameterization trick ([Kingma and Welling, 2014](#)).

During inference, latent variables obtained via prior networks and predicted question type qt' are fed to the question decoder, which corresponds to red dashed arrows in [Figure 2](#). The inference process is as follows:

- (1) Sample triple-level LV: $\mathbf{z}^t \sim q_\phi(\mathbf{z}^t|\mathbf{p})$ ¹.
- (2) Sample answer LV: $\mathbf{z}^a \sim p_\varphi(\mathbf{z}^a|\mathbf{p}, \mathbf{z}^t)$.
- (3) Sample question LV: $\mathbf{z}^q \sim p_\varphi(\mathbf{z}^q|\mathbf{p}, \mathbf{z}^t, \mathbf{z}^a)$.
- (4) Predict question type: $qt \sim p_\theta(qt|\mathbf{z}^q, \mathbf{z}^t, \mathbf{p})$.
- (5) Generate question: $\mathbf{q} \sim p_\theta(\mathbf{z}^q, \mathbf{z}^t, \mathbf{p}, qt)$.

3 Experiments

In this section, we conduct experiments to evaluate our proposed method. We first introduce some empirical settings, including dataset, hyper-parameters, baselines, and evaluation measures. Then we illustrate our results under both automatic and human evaluations. Finally, we give out some cases generated by different models and do further analyses over our method.

3.1 Dataset

We apply our model on a large-scale CQG corpus² extracted from Reddit³ by Wang et al. (2019). There are over 1.2 million PQA triples which have been divided into training/validation/test set with the number of 1,164,345/30,000/30,000. The dataset has been tokenized into words using the NLTK tokenizer (Bird et al., 2009). The average number of words in post/question/answer is 18.84/19.03/19.30, respectively. Following Fan et al. (2018) and Wang et al. (2019), we categorize questions in training and validation set into 9 types based on interrogative words, i.e., “what”, “when”, “where”, “who”, “why”, “how”, “can (could)”, “do (did, does)”, “is (am, are, was, were)”

3.2 Hyper-parameter Settings

We keep the top 40,000 frequent words as the vocabulary and the sentence padding length is set to 30. The dimension of GRU layer, word embedding and latent variables is 300, 300, and 100. The prior networks and MLPs have one hidden layer with size 300 and tanh non-linearity, while the number of hidden layers in recognition networks for both triple-level and utterance-level variables is 2. We apply dropout ratio of 0.2 during training. The mini-batch size is 64. For optimization, we use Adam (Kingma and Ba, 2015) with a learning rate of 1e-4. In order to alleviate degeneration problem of variational framework (Park et al., 2018), we

¹Inspired by Park et al. (2018), using \mathbf{z}^t inferred from post with the posterior distribution is better than sampling it from the prior one, i.e., a standard Gaussian distribution.

²https://drive.google.com/drive/folder/1wNG30YPHiMc_ZNyE3BH5waluVtR81lpG

³<http://www.reddit.com>

apply KL annealing, word drop (Bowman et al., 2016) and bag-of-words (BOW) loss (Zhao et al., 2017)⁴. The KL multiplier λ gradually increases from 0 to 1, and the word drop probability is 0.25. We use Pytorch to implement our model, and the model is trained on Titan Xp GPUs.

3.3 Baselines

We compare our methods with four groups of representative models: (1) **S2S-Attn**: A simple Seq2Seq model with attention mechanism (Shang et al., 2015). (2) **CVAE&kgCVAE**: The CVAE model integrates an extra BOW loss to generate diverse questions. The kgCVAE is a knowledge-guided CVAE that utilizes some linguistic cues (question types in our experiments) to learn meaningful latent variables (Zhao et al., 2017). (3) **STD&HTD**: The STD uses soft typed decoder that estimates a type distribution over word types, and the HTD uses hard typed decoder that specifies the type of each word explicitly with Gumbel-softmax (Wang et al., 2018). (4) **RL-CVAE**: A reinforcement learning method that regards the coherence score (computed by a one-to-one matching network) of a pair of generated question and answer as the reward function (Wang et al., 2019). RL-CVAE is the first work to utilize the future information, i.e., answer, and is also the state-of-the-art model for CQG⁵.

Additionally, we also conduct ablation study to better analyze our method as follows: (5) **GTM- \mathbf{z}^t** : GTM without the triple-level latent variable, which means \mathbf{z}^t is not included in the prior and posterior distributions of both \mathbf{z}^p and \mathbf{z}^a . (6) **GTM-a**: the variant of GTM that does not take answer into account. That is, answer decoder and \mathbf{z}^a are removed from the loss function and the prior and posterior distributions of \mathbf{z}^q . Besides, \mathbf{z}^t here does not capture the semantics from answer. (7) **GTM- $\mathbf{z}^q/\mathbf{z}^a$** : GTM variant in which distributions of \mathbf{z}^q are not conditioned on \mathbf{z}^a , i.e., the fact that the content of question is also controlled by answer is not modelled explicitly by latent variables.

In our model, we use an MLP to predict question types during inference, which is different from the conditional training (CT) methods (Li et al., 2016b; Zhou et al., 2018; Shen and Feng, 2020)

⁴The total BOW loss is calculated as the sum of all BOW losses between each latent variable and q/a . Please refer to Park et al. (2018) for more details.

⁵For those methods with open-source codes, we run the original codes; otherwise, we re-implement them based on the corresponding paper.

Model	Embedding Metrics			Diversity		BLEU Scores		RUBER Scores	
	Average	Extrema	Greedy	Dist-1	Dist-2	BLEU-1	BLEU-2	RubG	RubA
S2S-Attn	0.634	0.322	0.413	0.0132	0.0830	0.0936	0.0298	0.584	0.622
CVAE	0.646	0.337	0.421	0.0160	0.1599	0.1422	0.0306	0.649	0.687
kgCVAE	0.647	0.332	0.425	0.0153	0.1587	0.1491	0.0310	0.650	0.682
STD	0.637	0.326	0.418	0.0144	0.1325	0.1327	0.0302	0.633	0.663
HTD	0.648	0.330	0.423	0.0154	0.1582	0.1475	0.0314	0.653	0.689
RL-CVAE	0.662	0.343	0.437	0.0161	0.1785	0.1503	0.0320	0.660	0.701
GTM- z^t	0.672	0.351	0.448	0.0165	0.1872	0.1521	0.0332	0.661	0.710
GTM-a	0.653	0.338	0.428	0.0158	0.1679	0.1482	0.0317	0.657	0.692
GTM- z^a/z^a	0.687	0.360	0.449	0.0170	0.1934	0.1528	0.0329	0.669	0.713
GTM	0.697	0.365	0.454	0.0176	0.2028	0.1537	0.0331	0.671	0.720

Table 2: Automatic evaluation results for different models based on four types of metrics.

that provide the controllable feature, i.e., question types, in advance for inference. Therefore, we do not consider CT-based models as comparable ones.

3.4 Evaluation Measures

To better evaluate our results, we use both quantitative metrics and human judgements in our experiments.

Automatic Metrics

For automatic evaluation, we mainly choose four kinds of metrics: (1) **BLEU Scores:** BLEU (Papineni et al., 2002) calculates the n-gram overlap score of generated questions against ground-truth questions. We use BLEU-1 and BLEU-2 here and normalize them to 0 to 1 scale. (2) **Embedding Metrics:** Average, Greedy and Extrema metrics are embedding-based and measure the semantic similarity between the words in generated questions and ground-truth questions (Serban et al., 2017; Liu et al., 2016). We use word2vec embeddings trained on the Google News Corpus⁶ in this part. Please refer to Serban et al. (2017) for more details. (3) **Dist-1& Dist-2:** Following the work of Li et al. (2016a), we apply *Distinct* to report the degree of diversity. *Dist-1/2* is defined as the ratio of unique uni/bi-grams over all uni/bi-grams in generated questions. (4) **RUBER Scores:** Referenced metric and Unreferenced metric Blended Evaluation Routine (Tao et al., 2018) has shown a high correlation with human annotation in open-domain conversation evaluation. There are two versions, one is RubG based on geometric averaging and the other is RubA based on arithmetic averaging.

Embedding metrics and BLEU scores are used to measure the similarity between generated and ground-truth questions. RubG/A reflects the se-

mantic coherence of PQ pairs (Wang et al., 2019), while Dist-1/2 evaluates the diversity of questions.

Human Evaluation Settings

Inspired by Wang et al. (2019), Shen et al. (2019), and Wang et al. (2018), we use following three criteria for human evaluation: (1) **Fluency** measures whether the generated question is reasonable in logic and grammatically correct. (2) **Coherence** denotes whether the generated question is semantically consistent with the given post. Incoherent questions include dull cases. (3) **Willingness** measures whether a user is willing to answer the question. This criterion is to justify how likely the generated questions can elicit further interactions.

We randomly sample 500 examples from test set, and generate questions using models mentioned above. Then, we send each post and corresponding 10 generated responses to three human annotators without order, and require them to evaluate whether each question satisfies criteria defined above. All annotators are postgraduate students and not involved in other parts of our experiments.

3.5 Experimental Results

Now we demonstrate our experimental results on both automatic evaluation and human evaluation.

Automatic Evaluation Results

Now we demonstrate our experimental results on both automatic evaluation and human evaluation. The automatic results are shown in Table 2. The top part is the results of all baseline models, and we can see that GTM outperforms other methods on all metrics (significance tests (Koehn, 2004), p -value < 0.05), which indicates that our proposed model can improve the overall quality of generated questions. Specifically, Dist-2 and RubA have been improved by 2.43% and 1.90%, respectively, compared to the state-of-the-art RL-CVAE model.

⁶<https://code.google.com/archive/p/word2vec/>

First, higher embedding metrics and BLEU scores show that questions generated by our model are similar to ground truths in both topics and contents. Second, taking answer into account and using it to decide the expression of question can improve the consistency of PQ pairs evaluated by RUBER scores. Third, higher distinct values illustrate that one-to-many mappings in PQ and QA pairs make the generated responses more diverse.

The bottom part of Table 2 shows the results of our ablation study, which demonstrates that taking advantage of answer information, modeling the shared background in entire triple, and considering one-to-many mappings in both PQ and QA pairs can help enhance the performance of our hierarchical variational model in terms of relevance, coherence and diversity.

Human Evaluation Results

As shown in Table 3, GTM can alleviate the problem of generating dull and deviated questions compared with other models (significance tests (Koehn, 2004), p -value < 0.05). Both our proposed model and the state-of-the-art model RL-CVAE utilize the answer information and the results of them could prove that answers assist the question generation process. Besides, GTM can produce more relevant and intriguing questions, which indicates the effectiveness of modeling the shared background and one-to-many mappings in CQG task. The inter-annotator agreement is calculated with the Fleiss’ kappa (Fleiss and Cohen, 1973). Fleiss’ kappa for Fluency, Coherence and Willingness is 0.493, 0.446 and 0.512, respectively, indicating “Moderate Agreement” for all three criteria.

3.6 Question-Answer Coherence Evaluation

Automatic metrics in Section “Automatic Metrics” are designed to compare generated questions with ground-truth ones (RUBER also takes the post information into consideration), but ignore answers in the evaluation process. To measure the semantic coherence between generated questions and answers, we apply two methods (Wang et al., 2019): (1) Cosine Similarity: We use the pre-trained Inference model⁷ (Conneau et al., 2017) to obtain sentence embeddings and calculate cosine similarity between the embeddings of generated responses

⁷The Inference model is trained to predict the meaning of sentences based on natural language inference, and the cosine similarity computed with it is more consistent with human’s judgements, which performs better than the pre-trained Transformer/BERT model in our experiments.

and answers. (2) Matching Score: We use the GRU-MatchPyramid (Wang et al., 2019) model that adds the MatchPyramid network (Pang et al., 2016) on top of a bidirectional GRU to calculate the semantic coherence. As shown in Table 4, questions generated by GTM are more coherent to answers. Attributing to the design of triple-level latent variable that captures the shared background, one-to-many

Model	Fluency	Coherence	Willingness
S2S-Attn	0.482	0.216	0.186
CVAE	0.462	0.484	0.428
kgCVAE	0.474	0.536	0.476
STD	0.488	0.356	0.286
HTD	0.526	0.504	0.414
RL-CVAE	0.534	0.578	0.508
GTM- \mathbf{z}^t	0.538	0.580	0.516
GTM-a	0.532	0.570	0.512
GTM- $\mathbf{z}^q/\mathbf{z}^a$	0.542	0.586	0.520
GTM	0.548	0.608	0.526

Table 3: Results for human evaluation.

Model	Cosine Similarity	Matching Score
S2S-Attn	0.498	5.306
CVAE	0.564	8.047
kgCVAE	0.578	8.054
STD	0.542	6.879
HTD	0.583	8.059
RL-CVAE	0.607	8.423
GTM- \mathbf{z}^t	0.613	8.427
GTM-a	0.605	8.424
GTM- $\mathbf{z}^q/\mathbf{z}^a$	0.618	8.472
GTM	0.629	8.517

Table 4: Evaluation results for QA coherence.

mappings in PQ and QA pairs, and relationship modeling for \mathbf{z}^q and \mathbf{z}^a , GTM can improve the relevance in QA pairs.

3.7 Case Study

In Table 5, we list the generated results of two posts from the test set to compare the performance of different models.

In the first case, both the post and answer mention two topics, “donation” and “song”, so the question is better to consider their relations. Besides, the answer here begins with “because”, then “why” and “what (reason)” questions are reasonable. For the second case, the post only talks about “pen”, while the answer refers to “ink”, which means there is a topic transition the question needs to cover. The second case shows the effectiveness of an answer that not only decides the expression of question but also improves the entire coherence of a tripe. Questions generated by GTM are more relevant to

Post	Question	Answer
and it's a donation to food banks, even if you dislike the song.	S2S-Attn: what does that mean? CVAE: what's the best way to do that? kgCVAE: is there anything else to buy? STD: is it a good donation? HTD: where are those food banks? RL-CVAE: why do we need to do the stupid thing? GTM: why not just donate money to food banks?	because i like the aspect of song criticising may reaching to the charts as well as the fact it goes to charity.
it's a very blue pen. i'm more into brighter colours, but this is still pretty nice.	S2S-Attn: what colors do you have? CVAE: are you colorblind? kgCVAE: it has nothing to complain. where's the pen? STD: what color are you using? HTD: what's the colour scheme? RL-CVAE: what kind of ink should i buy? GTM: i'll take it though. do you also sell the ink?	i have some coming, hopefully it will be here early next week. right now it's got green ink and it's really messing me up.

Table 5: Two cases comparison among GTM and other baselines.

both posts and answers, and could attract people to give an answer to them. However, other baselines may generate dull or deviated responses, even the RL-CVAE model that considers the answer information would only contain the topic words in answers (e.g., the question in case two), but fail to ensure the PQA coherence.

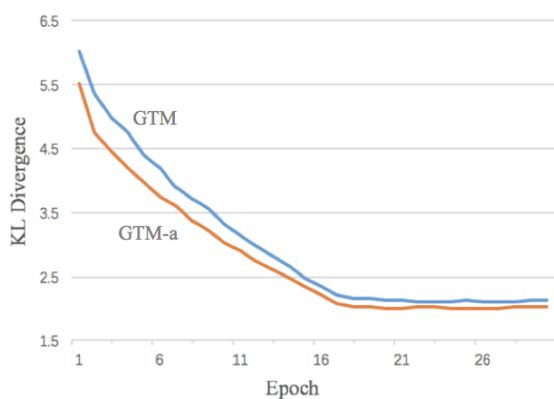


Figure 3: Total KL divergence (per word) of all latent variables in GTM and GTM-a model (first 30 epochs of validation set).

3.8 Further Analysis of GTM

Variational models suffer from the notorious degeneration problem, where the decoders ignore latent variables and reduce to vanilla Seq2Seq models (Zhao et al., 2017; Park et al., 2018; Wang et al., 2019). Generally, KL divergence measures the amount of information encoded in a latent variable. In the extreme case where the KL divergence of latent variable z equals to zero, the model completely ignores z , i.e., it degenerates. Figure 3 shows that the total KL divergence of GTM model maintains around 2 after 18 epochs indicating that the degen-

eration problem does not exist in our model and latent variables can play their corresponding roles.

4 Related Work

The researches on open-domain dialogue systems have developed rapidly (Majumder et al., 2020; Zhan et al., 2021; Shen et al., 2021), and our work mainly touches two fields: open-domain conversational question generation (CQG), and context modeling in dialogue systems. We introduce these two fields as follows and point out the main differences between our method and previous ones.

4.1 CQG

Traditional question generation (TQG) has been widely studied and can be seen in reading comprehension (Zhou et al., 2019; Kim et al., 2019), sentence transformation (Vanderwende, 2008), question answering (Li et al., 2019; Nema et al., 2019), visual question generation (Fan et al., 2018) and task-oriented dialogues (Li et al., 2017). In such tasks, finding information via a generated question is the major goal and the answer is usually part of the input. Different from TQG, CQG aims to enhance the interactivenss and persistence of conversations (Wang et al., 2018). Meanwhile, the answer is the “future” information which means it is unavailable in the inference process. Wang et al. (2018) first studied on CQG, and they used soft and hard typed decoders to capture the distribution of different word types in a question. Hu et al. (2018) added a target aspect in the input and proposed an extended Seq2Seq model to generate aspect-specific questions. Wang et al. (2019) devised two methods based on either reinforcement learning or generative adversarial network (GAN)

to further enhance semantic coherence between posts and questions under the guidance of answers.

4.2 Context Modeling in Dialogue Systems

Existing methods mainly focus on the historical context in multi-turn conversations, and hierarchical models occupy a vital position in this field. Serban et al. (2016) proposed the hierarchical recurrent encoder-decoder (HRED) model with a context RNN to integrate historical information from utterance RNNs. To capture utterance-level variations, Serban et al. (2017) raised a new model Variational HRED (VHRED) that augments HRED with CVAEs. After that, VHCR (Park et al., 2018) added a conversation-level latent variable on top of the VHRED, while CSRR (Shen et al., 2019) used three-hierarchy latent variables to model the complex dependency among utterances. In order to detect relative utterances in context, Tian et al. (2017) and Zhang et al. (2018) applied cosine similarity and attention mechanism, respectively. HRAN (Xing et al., 2018) combined the attention results on both word-level and utterance-level. Besides, the future information has also been considered for context modeling. Shen et al. (2018) separated the context into history and future parts, and assumed that each of them conditioned on a latent variable is under a Gaussian distribution. Feng et al. (2020) used future utterances in the discriminator of a GAN, which is similar to Wang et al. (2019).

The differences between our method and aforementioned ones in Section 4.1 and 4.2 are: (1) Rather than dividing PQA triples into two parts, i.e., PQ (history and current utterances) and QA (current and future utterances) pairs, we model the entire coherence by utilizing a latent variable to capture the share background in a triple. (2) Instead of regarding the relationship between question and answer as a text matching task that lacks the consideration of diversity, we incorporate utterance-level latent variables to help model one-to-many mappings in both PQ and QA pairs.

5 Conclusion

We propose a generative triple-wise model for generating appropriate questions in open-domain conversations, named GTM. GTM models the entire background in a triple and one-to-many mappings in PQ and QA pairs simultaneously with latent variables in three hierarchies. It is trained in a one-stage end-to-end framework without pre-training

like the previous state-of-the-art model that also takes answer into consideration. Experimental results on a large-scale CQG dataset show that GTM can generate fluent, coherent, informative as well as intriguing questions.

Acknowledgements

We would like to thank all the reviewers for their insightful and valuable comments and suggestions.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. O’Reilly Media, Inc.”.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Chaotao Chen, Jinhua Peng, Fan Wang, Jun Xu, and Hua Wu. 2019. Generating multiple diverse responses with multi-mapping and posterior mapping selection. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 4918–4924. AAAI Press.
- Kyunghyun Cho, Bart van Merriënboer Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. *Supervised learning of universal sentence representations from natural language inference data*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhihao Fan, Zhongyu Wei, Piji Li, Yanyan Lan, and Xuanjing Huang. 2018. A question type driven framework to diversify visual question generation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4048–4054. AAAI Press.

- Shaoxiong Feng, Hongshen Chen, Kan Li, and Dawei Yin. 2020. Posterior-gan: Towards informative and coherent response generation with posterior generative adversarial network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7708–7715.
- Joseph L Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Xiang Gao, Sungjin Lee, Yizhe Zhang, Chris Brockett, Michel Galley, Jianfeng Gao, and Bill Dolan. 2019. [Jointly optimizing diversity and relevance in neural response generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1229–1238, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wenpeng Hu, Bing Liu, Jinwen Ma, Dongyan Zhao, and Rui Yan. 2018. Aspect-based question generation.
- Yanghoon Kim, Hwanhee Lee, Joongbo Shin, and Kyomin Jung. 2019. Improving neural question generation using answer separation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6602–6609.
- Diederick P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *the 3rd International Conference on Learning Representations*.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *the 2nd International Conference on Learning Representations*.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Dong Bok Lee, Seanie Lee, Woo Tae Jeong, Donghwan Kim, and Sung Ju Hwang. 2020. [Generating diverse and consistent QA pairs from contexts with information-maximizing hierarchical conditional VAEs](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 208–224, Online. Association for Computational Linguistics.
- Jingjing Li, Yifan Gao, Lidong Bing, Irwin King, and Michael R. Lyu. 2019. [Improving question generation with to the point context](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3216–3226, Hong Kong, China. Association for Computational Linguistics.
- 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, San Diego, California. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016b. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.
- Jiwei Li, Alexander H Miller, Sumit Chopra, Marc’Aurelio Ranzato, and Jason Weston. 2017. Learning through dialogue interactions by asking questions. *ICLR*.
- 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, Austin, Texas. Association for Computational Linguistics.
- Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian McAuley. 2020. [Like hiking? you probably enjoy nature: Persona-grounded dialog with commonsense expansions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9194–9206, Online. Association for Computational Linguistics.
- Preksha Nema, Akash Kumar Mohankumar, Mitesh M. Khapra, Balaji Vasan Srinivasan, and Balaraman Ravindran. 2019. [Let’s ask again: Refine network for automatic question generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3314–3323, Hong Kong, China. Association for Computational Linguistics.
- Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- 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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

- Yookoon Park, Jaemin Cho, and Gunhee Kim. 2018. [A hierarchical latent structure for variational conversation modeling](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1792–1801, New Orleans, Louisiana. Association for Computational Linguistics.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 3776–3783.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 3295–3301.
- 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: Long Papers)*, pages 1577–1586, Beijing, China. Association for Computational Linguistics.
- Lei Shen and Yang Feng. 2020. [CDL: Curriculum dual learning for emotion-controllable response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 556–566, Online. Association for Computational Linguistics.
- Lei Shen, Yang Feng, and Haolan Zhan. 2019. [Modeling semantic relationship in multi-turn conversations with hierarchical latent variables](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5497–5502, Florence, Italy. Association for Computational Linguistics.
- Lei Shen, Haolan Zhan, Xin Shen, and Yang Feng. 2021. Learning to select context in a hierarchical and global perspective for open-domain dialogue generation. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7438–7442. IEEE.
- Xiaoyu Shen, Hui Su, Wenjie Li, and Dietrich Klakow. 2018. [NEXUS network: Connecting the preceding and the following in dialogue generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4316–4327, Brussels, Belgium. Association for Computational Linguistics.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. Ruber: An unsupervised method for automatic evaluation of open-domain dialog systems. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Zhiliang Tian, Rui Yan, Lili Mou, Yiping Song, Yansong Feng, and Dongyan Zhao. 2017. [How to make context more useful? an empirical study on context-aware neural conversational models](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 231–236, Vancouver, Canada. Association for Computational Linguistics.
- Lucy Vanderwende. 2008. The importance of being important: Question generation. In *Proceedings of the 1st Workshop on the Question Generation Shared Task Evaluation Challenge*, Arlington, VA.
- Weichao Wang, Shi Feng, Daling Wang, and Yifei Zhang. 2019. [Answer-guided and semantic coherent question generation in open-domain conversation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5066–5076, Hong Kong, China. Association for Computational Linguistics.
- Yansen Wang, Chenyi Liu, Minlie Huang, and Liqiang Nie. 2018. [Learning to ask questions in open-domain conversational systems with typed decoders](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2193–2203, Melbourne, Australia. Association for Computational Linguistics.
- Chen Xing, Yu Wu, Wei Wu, Yalou Huang, and Ming Zhou. 2018. Hierarchical recurrent attention network for response generation. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5610–5617.
- Zhou Yu, Ziyu Xu, Alan W Black, and Alexander Rudnicky. 2016. Strategy and policy learning for non-task-oriented conversational systems. In *Proceedings of the 17th annual meeting of the special interest group on discourse and dialogue*, pages 404–412.
- Haolan Zhan, Hainan Zhang, Hongshen Chen, Zhuoye Ding, Yongjun Bao, and Yanyan Lan. 2021. [Augmenting knowledge-grounded conversations with sequential knowledge transition](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5621–5630, Online. Association for Computational Linguistics.
- Weinan Zhang, Yiming Cui, Yifa Wang, Qingfu Zhu, Lingzhi Li, Lianqiang Zhou, and Ting Liu. 2018. Context-sensitive generation of open-domain conversational responses. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2437–2447.
- 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: Long Papers*), pages 654–664, Vancouver, Canada. Association for Computational Linguistics.

Hao Zhou, Minlie Huang, Tianyang Zhang, Xiaoyan Zhu, and Bing Liu. 2018. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Wenjie Zhou, Minghua Zhang, and Yunfang Wu. 2019. [Question-type driven question generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6032–6037, Hong Kong, China. Association for Computational Linguistics.