
HSCJN: A Holistic Semantic Constraint Joint Network for Diverse Response Generation

Yiru Wang **Pengda Si** **Zeyang Lei** *
Tsinghua Shenzhen International Graduate School Baidu Inc.
Tsinghua University, P. R. China Beijing, P. R. China
{wangyiru17, spd18}@mails.tsinghua.edu.cn zeyanglei@gmail.com

Guangxu Xun **Yujiu Yang** †
Department of Computer Science Tsinghua Shenzhen International Graduate School
University of Virginia, USA Tsinghua University, Shenzhen, P. R. China
gx5bt@virginia.edu yang.yujiu@sz.tsinghua.edu.cn

Abstract

The sequence-to-sequence (Seq2Seq) model generates target words iteratively given the previously observed words during decoding process, which results in the loss of the holistic semantics in the target response and the complete semantic relationship between responses and dialogue histories. In this paper, we propose a generic diversity-promoting joint network, called Holistic Semantic Constraint Joint Network (HSCJN), enhancing the global sentence information, and then regularizing the objective function with penalizing the low entropy output. Our network introduces more target information to improve diversity, and captures direct semantic information to better constrain the relevance simultaneously. Moreover, the proposed method can be easily applied to any Seq2Seq structure. Extensive experiments on several dialogue corpuses show that our method effectively improves both semantic consistency and diversity of generated responses, and achieves better performance than other competitive methods.

1 Introduction

Recently, dialogue systems have attracted increasing attention in both academia and industry because of their potential applications and commercial values. Sequence-to-sequence (Seq2Seq) models form the cornerstone of popular response generation models (Serban et al., 2016; Sordoni et al., 2015; Shang et al., 2015). However, neural dialogue systems based on Seq2Seq models tend to repeatedly generate universal and boring responses like “I don’t know.”, “Thank you.”. Although being widely applied, conventional Maximum Likelihood Estimation (MLE) training could cause the low-diversity problem above (Li et al., 2016b). Since high frequency words make up a big proportion of the training set, MLE encourages the model to excessively generate high frequency words.

Moreover, when training a Seq2Seq model traditionally, we iteratively maximize the log predictive likelihood of each true token in the target sequence, given the previously decoded tokens. Therefore, the model can only see the previous information during learning, unable to grasp the holistic information of the target sequence when decoding tokens. This also leads to the loss of complete semantic relationship between target sequences and source sequences.

*This work was done when Zeyang Lei was in Tsinghua University.

†Corresponding author. This work was supported partially by Guangdong Basic and Applied Basic Research Foundation (No.2019A1515011387), and Shenzhen Basic Research Foundation (No.JCYJ201704is170118573).

As discussed, we argue that the current learning strategy heavily limits the Seq2Seq models to generate highly diverse responses, and the holistic semantic information of the target response as well as the global semantic relationship between responses and dialog histories are missing during generation process. Most previous solutions simply rely on external information or post-processing models to mask the deficiencies of the Seq2Seq model, while the problem of the Seq2Seq model itself has not been addressed. Therefore, in this paper, we hope to fully exploit the learning potential of the Seq2Seq models without any external information to improve diversity while better to constrain the semantic relevance of generated responses simultaneously. We propose a Holistic Semantic Constraint Joint Network (HSCJN) to predict the subsequent word set in each target utterance for direct supervision in decoding, which directly introduces more linguistic information from target utterances to increase diversity. More specifically, in the HSCJN, we require each hidden state in the decoder to predict the words in the target utterance which remain ungenerated, and the initial state of the decoder is required to predict all the words in the target utterance. Since the HSCJN enables the decoder network to see all words in the target utterance at every time step, our model also more likely captures direct semantic information such as keywords in target utterances to enhance relevance.

In this way, the relationship of representation spaces between source sequences and target sequences, and the transition between different decoder states could be better constrained. In addition, we consider that the entropy of the output distribution is low if the model is over-confident about high frequency words. Penalizing the low entropy output distribution can help regularize the model, optimize the predicted output distribution and alleviate the over-estimation of high frequency words. Therefore, we devise a maximum entropy based regularizer. Our learning framework can be used as a general joint training method with Seq2Seq models and requires no additional data or annotation. In general, our contributions are summarized as follows:

- We devise a joint training network to introduce future information in the decoding stage in open-domain dialogue generation, which can be applied to any Seq2Seq neural model. Our network introduces more linguistic information from target utterances to increase diversity, and likely captures key semantic information such as keywords in target utterances to enhance relevance in a direct manner.
- We regularize the model by penalizing low entropy output distribution at each time step in the decoder to alleviate the over-estimation of high frequency words, which also enables the loss function to consider every word in the vocabulary to improve diversity.
- The experimental results on multi-turn dialogue datasets show the effectiveness of our method in terms of both diversity and relevance of generated responses.

2 Related Work

The diversity of generated responses is an important issue of common concern. Serban et al. (2017) and Zhao et al. (2017) proposed to introduce variational auto-encoders (VAEs) to Seq2Seq models to bring informativeness by increasing variability. Some researches proposed several beam search based approaches (Li et al., 2017; Song et al., 2018; Vijayakumar et al., 2016). However, this kind of methods merely provide a criterion for reweighing response candidates, rather than producing more diverse responses in the first place. Besides, another previous works introduce additional information or knowledge such as the context (Serban et al., 2016; Tian et al., 2017; Yao et al., 2017), keyword (Serban et al., 2016; Xing et al., 2017; Yao et al., 2017) or knowledge-base (Young et al., 2018; Ghazvininejad et al., 2018) into the response generation process to produce informative content. Although being effective, these approaches actually bypass the low-diversity problem by introducing the randomness of stochastic latent variables or additional information. The underlying Seq2Seq model remains sub-optimal in terms of diversity. Li et al. (2016a) proposed to use a Maximum Mutual Information (MMI) as an optimization objective to maximize the mutual information between messages and responses, but the MMI objective is used only during test time, and relies on many extra modules, like reverse models and beam search. Zhang et al. (2018) proposed the Adversarial Information Maximization (AIM) model which considers explicitly maximizing mutual information during training to generate informative responses, but it still needs to train an extra backward model generating source from target, and implemented with complicated adversarial training strategy.

In other tasks, the word prediction technique has been applied in the neural machine translation (Weng et al., 2017; L’Hostis et al., 2016). Lin et al. (2019) adds entropy to the loss function to make the sparse distribution more specifically concentrate on a small set of video segments in the VQA task. Unlike we consider entropy on the entire vocabulary, it only considers the entropy of audio, video, and words in the current sentence on the video segment.

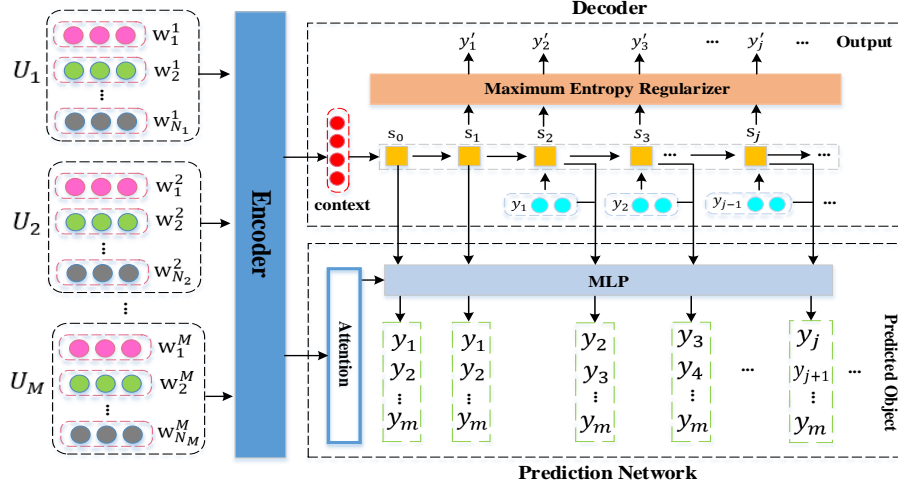


Figure 1: The architecture of the proposed HSCJN for multi-turn dialogue generation. The prediction network only works during the training phase which helps improve the Seq2Seq model for inference.

3 HSCJN Model

3.1 Task Definition

Given a dialogue as a sequence of M utterances $\mathbb{U} = \{U_1, \dots, U_M\}$, and each U_i as a sequence of N_i tokens $U_i = \{w_{i1}, \dots, w_{iN_i}\}$, where w_{ij} represents the token at position j in utterance i from the vocabulary V , our task is to generate a response $Y = \{y_1, y_2, \dots, y_m\}$ that is not only fluent and grammatical but also not repeated and trivial in content. Essentially, the goal is to estimate the conditional probability:

$$P(Y|\mathbb{U}) = \prod_{t=1}^m P(y_t|\mathbb{U}, y_{<t}), \text{ with } \mathbb{U} = \left\{ [w_{1k}]_{k=1}^{N_1}, \dots, [w_{Mk}]_{k=1}^{N_M} \right\} \quad (1)$$

3.2 An Overview of HSCJN

Figure 1 demonstrates the architecture of our model, in which we join a prediction network with the decoder network on each hidden state of the decoder. The encoder takes the word embedding sequences of context utterances as inputs, and obtains the hidden representations of the context. The decoder starts the generation of the target sequence from the initial state s_0 . Since the initial state is responsible for the generation of the whole target sequence, we optimize the initial state by making prediction for all the target words to contain comprehensive target information. Similarly, at each time step in the decoder, we introduce the prediction network to predict the word set of the target subsequence that has not yet been generated. The response is generated from the decoder, HSCJN applies a constraint network to each hidden state in the decoder to introduce more direct language information from the Seq2Seq model. There is a specific objective function for the HSCJN. In addition, we optimize the output distribution at each decoding step by adding a maximum entropy based regularizer to the final objective function to generate the predictive response.

3.3 Holistic Semantic Constraint Joint Network Design

In HSCJN, we require the hidden state at each time step in the decoder to predict the word set containing target words which remain ungenerated in the target utterance, where the order of words is not considered and we assume target words are independent with each other. In this way, at each time step, the decoder generates words not only conditioned on the previously generated subsequences within the original decoder network, but also under consideration of the future words not yet seen in the target sequence through our HSCJN. That is, our joint network HSCJN can introduce and utilize the global sentence information in target utterances for every token generation, beneficial for both diversity and relevance. Specifically, for each time step j in the decoder, the hidden state s_j is

required to predict the word collection of $Y_{j\sim m} \triangleq (y_j, y_{j+1}, \dots, y_m)$. The conditional probability P_j of the prediction task in the HSCJN at hidden state s_j is defined as follows:

$$P_j(Y_{j\sim m}|\Psi) = \prod_{t=j}^m P(y_t|\Psi), \text{ with } P(y_t|\Psi) = MLP([e(y_{j-1}); s_j; c_j]) \quad (2)$$

where, $\Psi \triangleq \{y_{<j}, U_1, \dots, U_M\}$, and the set $Y_{j\sim m}$ is the word set of the future subsequence in target response Y at time step j . MLP is a multi-layer perceptron with two hidden layers using $\tanh(\cdot)$ as an activation function, followed by one output layer with $\text{sigmoid}(\cdot)$ acting on each neuron. Here, we predict the target word set in a multi-classification way. $e(y_{j-1})$ is the embedding of the word y_{j-1} , and c_j is the context vector from the attention mechanism(Luong et al., 2015).

$$c_j = \sum_{i=1}^T a_{ji} h_i, \text{ with } a_{ji} = \frac{\exp(\tanh(W_c[s_{j-1}, h_i]))}{\sum_{k=1}^T \exp(\tanh(W_c[s_{j-1}, h_k]))} \quad (3)$$

where W_c is weight parameters, T is the input sequence length, h_i is the hidden state of the encoder RNN at time step i , and s_{j-1} is the hidden state of the decoder RNN at the previous time step $j - 1$.

Specifically, for the initial state s_0 , HSCJN requires it to predict the word set containing all target words, so as to compress the overall information of the target sequence into the initial state. Therefore, the decoder can see the entire target sequence at the initial time step through the HSCJN. The conditional probability P_0 of the HSCJN at the initial state s_0 is defined as follows:

$$P_0(\tilde{Y}|\mathbb{U}) = \prod_{t=1}^m P_0(y_t|\mathbb{U}), \text{ with } P_0(y_t|\mathbb{U}) = MLP([s_0; c_0]) \quad (4)$$

where \tilde{Y} is the word set containing all target words in the target response Y , and c_0 is the context vector from the attention mechanism for the initial state s_0 .

To optimize the HSCJN network, we add an extra likelihood function L_{WLP} into the training procedure:

$$L_{WLP} = -\frac{1}{m} \log P_0 - \sum_{j=1}^m \frac{1}{m-j+1} \log P_j \quad (5)$$

where P_0 and P_j are as previously defined, the coefficient of the logarithm is used to calculate the average probability of each prediction. This loss function is used to guide HSCJN to accurately introduce the expected target semantic information.

3.4 Output Distribution Regularizer

When a dialogue model generates universal responses, the prediction of high frequency words is too confident, that is, the entire output probability distribution is concentrated on high frequency words. In result, the entropy of the output distribution is low. We consider that maximizing the entropy of the output distribution at each decoding step could help to regularize the model and produce more diverse responses. By this means, the token-level distribution $P(y_t|\mathbb{U}, y_{<t})$ is better constrained to relieve over-estimation of high frequency words. Therefore, we add a negative entropy to the negative log-likelihood loss function during training. To minimize the overall loss function, the model encourages the maximization of entropy. Specifically, the loss is expressed as:

$$L_{ME} = -\sum_{t=1}^m H(p(y_t|\tilde{\Psi})), \text{ with } H(p(y_t|\tilde{\Psi})) = -\sum_{i=1}^{|V|} p(w_i|\tilde{\Psi}) * \log p(w_i|\tilde{\Psi}) \quad (6)$$

Where $\tilde{\Psi} \triangleq \{y_{<t}, U_1, \dots, U_M\}$, $H(\cdot)$ is the entropy of the output distribution at decoding step t , $|V|$ is the length of the vocabulary V , and w_i represents a word in the vocabulary.

This loss function not only penalizes the low entropy output distribution when predicting each token, but also considers the entropy over the entire vocabulary, so that the model likely takes into account more words to increase diversity.

3.5 Loss Function

We add L_{WLP} and L_{ME} to the original negative log-likelihood loss function. The final loss function for model training is as follows:

$$L = -\log P(Y|\mathbb{U}) + \alpha L_{WLP} + \beta L_{ME} \quad (7)$$

Table 1: The results of 1-turn response generation on DailyDialog and OpenSubtitles datasets

DailyDialog Corpus							
Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Distinct-1	Distinct-2	Distinct-3
Attnseq2seq	7.17	0.15	0.03	0.02	0.031/247	0.092/523	0.162/754
HRED	9.44	1.85	0.83	0.36	0.049/276	0.145/516	0.252/651
VHRED	11.28	2.30	1.05	0.44	0.068/418	0.197/778	0.342/1007
HSCJN	11.05	2.60	1.29	0.56	0.075/463	0.242/954	0.424/1235
HSCJN(w/o ME)	11.48	2.55	1.28	0.69	0.084/536	0.276/1149	0.475/1481
HSCJN(w/o PN)	10.57	2.30	1.13	0.58	0.063/386	0.200/788	0.357/1043
OpenSubtitles Corpus							
Attnseq2seq	13.27	0.88	0.28	0.12	0.006/530	0.022/1623	0.043/2741
HRED	12.50	1.25	0.47	0.16	0.007/459	0.027/1391	0.050/2039
VHRED	14.64	1.28	0.36	0.13	0.007/707	0.034/2627	0.073/4855
HSCJN	13.29	1.45	0.55	0.26	0.012/883	0.063/3550	0.131/6056
HSCJN(w/o ME)	13.15	1.48	0.54	0.23	0.011/880	0.056/3411	0.114/5773
HSCJN(w/o PN)	12.61	1.44	0.51	0.25	0.008/621	0.039/2231	0.078/3650

Table 2: 2-turns dialog generation results on DailyDialog Corpus

Model	BLEU-1	BLEU-2	BLEU-3	BLEU-4	Distinct-1	Distinct-2	Distinct-3
HRED	10.23	1.78	0.72	0.33	0.037/437	0.107/827	0.192/1104
VHRED	12.24	2.26	0.91	0.37	0.052/663	0.150/1265	0.269/1749
HSCJN	11.63	2.25	1.01	0.54	0.053/689	0.175/1500	0.313/2056

where α and $\beta \in [0, 1]$ are weight coefficients, which control the strength of the joint prediction task and the output distribution regularizer respectively. Our HSCJN builds a training objective at sentence level instead of the traditional token-level transition, considering the complete linguistic information in target utterances for every token generation.

4 Experiment

4.1 Data Preparation

We evaluate our proposed method on two multi-turn dialogue datasets, DailyDialog (Li et al., 2017) and OpenSubtitles (Tiedemann, 2012).

DailyDialog: It is a high-quality and less noisy dataset, which contains 13,118 multi-turn dialogues, separated into training/validation/test sets with 11,118/1,000/1,000 conversations. For the computational efficiency, we remove the dialogues with more than 300 tokens, which only makes up a small proportion of the whole dataset, and finally our training/validation/test sets of DailyDialog dataset contain 10,712 / 976 / 960 conversations respectively.

OpenSubtitles: It is a collection of movie subtitles. Following previous work (Xu et al., 2018), we treat each turn in the dataset as the target text and the two previous sentences as the source text. We randomly sample 200,000 / 50,000 / 10,000 dialogues for training, validation, and testing, respectively. Similarly, we also remove dialogues with more than 300 tokens, and finally our training/validation/test sets of OpenSubtitles dataset contain 199,992/ 49,995 / 9,984 dialogues respectively.

4.2 Baselines

AttnSeq2Seq: A vanilla Seq2Seq model with attention mechanism (Bahdanau et al., 2014). The encoder and decoder are both recurrent neural networks (RNN) with LSTM as the basic cell, and the encoder RNN is bidirectional.

HRED: HRED (Serban et al., 2016) considers dialogue history in multi-turn dialogue generation at two levels: a sequence of words for each utterance and a sequence of utterances, and models this hierarchy of conversations accordingly.

VHRED: VHRED (Serban et al., 2017) augments the HRED model with a stochastic latent variable at the decoder, trained by maximizing a variational lower-bound on the log-likelihood. The latent variable helps facilitate the generation of long utterances with more information content.




Figure 2: Word frequency distribution on generated responses on OpenSubtitles dataset. We draw the ten most frequent words, excluding punctuations.

4.3 Model Settings

Our proposed method is generic since it can be combined with any Seq2Seq model. In our experiments, we use HRED as the basis of our learning framework. We initialize the recurrent parameter matrices as orthogonal matrices while all the bias vectors are set to $\mathbf{0}$. Other parameters are initialized by sampling from the Gaussian stochastic distribution $\mathcal{N}(0, 0.01)$. The vocabularies are limited to the most frequent 25K and 30K words for DailyDialog dataset and OpenSubtitles dataset respectively. We apply GRU with 500 hidden states and GRU with 1000 hidden states to the encoders at word-level and utterance-level respectively, and LSTM with 500 hidden states to the decoder. The dimension of word embedding is set to 300. We use the Adam optimizer (Kingma and Ba, 2014) to update the parameters, with a batch size of 8. The learning rate is 0.0002 and the dropout rate is 0.75. Meanwhile, we set α to 1 and β to 0.13. For decoding during test time, we simply decode until the end-of-utterance symbol *eou* occurs, using a beam search with a beam width of 5. All models in the baselines are implemented with the same settings.

4.4 Automatic Evaluation

We adopt BLEU (Chen and Cherry, 2014; Papineni et al., 2002), Distinct-1, Distinct-2 and Distinct-3 (Li et al., 2016a) to evaluate the models at the quality and diversity level. The higher BLEU values demonstrate the responses are closer to the ground truth. Distinct-1, Distinct-2 and Distinct-3 are the proportion and number of distinct unigrams, bigrams, and trigrams in all the generated tokens, respectively. Higher Distinct- n values are better for the overall diversity.

Table 1 shows the experimental results of 1-turn response generation on DailyDialog corpus and OpenSubtitles corpus. It is obvious that our HSCJN generates remarkably more distinct unigrams, bigrams and trigrams than all the baselines on both two datasets. Besides, our model achieves the highest BLEU-2/3/4 values on both two datasets, compared with all the baseline models. Confirmed by our experiments, our model achieves an excellent performance in terms of both quality and diversity, regardless of the scale of datasets.

Furthermore, we conduct experiments on 2-turns dialogue generation. Given dialogue histories as input, we require models to generate the next two consecutive utterances. Since dialogues in the OpenSubtitles dataset contain only three turns, we conduct 2-turns response generation on the DailyDialog corpus. The results of 2-turns dialogue generation by our model and the other two multi-turn dialogue generation models are showed in Table 2. From the results, our model exceeds all baseline models with diversity improvement in multi-turn dialogue generation, and even higher BLEU-3 and BLEU-4.

To verify the effectiveness of our model in optimizing the output distribution, we perform word segmentation and word frequency statistics on the generated responses. Figure 2 draws the distribution of ten most frequent words in responses generated by corresponding models on the OpenSubtitles dataset, excluding punctuations. “Natural” represents the natural distribution of the ground truth. The horizontal axis represents the rank of word frequency, and the vertical axis is word frequency values. The curves above columns fit the distribution trends of the word frequencies for different models, whose colors are consistent with the columns. It shows that the frequencies of words generated by our model are not as high as the baseline models, and the frequency distribution is flatter. Moreover,

Table 3: Manual evaluation results for 1-turn response generation

Score	Dailydialog Corpus				OpenSubtitles Corpus			
	0	+1	+2	Average	0	+1	+2	Average
AttnSeq2Seq	68.4%	22.2%	9.4%	0.410	47.9%	34.6%	17.5%	0.708
HRED	55.2%	30.7%	14.2%	0.590	28.6%	50.0%	21.4%	0.928
VHRED	47.7%	30.0%	22.3%	0.747	29.4%	39.8%	30.8%	1.014
HSCJN	46.2%	28.8%	25.0%	0.788	21.4%	45.2%	33.4%	1.120

Table 4: Manual evaluation results for 2-turns dialog generation on DailyDialog Corpus

Score	0	+1	+2	Average
HRED	54.0%	24.4%	21.6%	0.676
VHRED	64.8%	17.2%	18.0%	0.532
HSCJN	53.2%	22.4%	24.4%	0.712

the distribution trends of our model and VHRED are basically consistent with the natural distribution, and our model is closer to the natural distribution in word frequency values than VHRED.

4.5 Ablation Study

We conduct the ablation study to examine the effectiveness of each mechanism, and the results are shown in Table 1. HSCJN(w/o ME) and HSCJN (w/o PN) represent the trained models without the maximum entropy regularizer and without the prediction sub-network respectively. On the DailyDialog dataset, HSCJN (w/o ME) generates the most distinct unigrams, bigrams and trigrams among all compared models, and also surpasses all the baseline models on Distinct- n metrics on the OpenSubtitles dataset, indicating that our joint network can obviously generate more diverse responses. HSCJN(w/o ME) also achieves higher BLEU scores, which proves that our joint network can directly capture semantic information to enhance relevance.

Since HSCJN(w/o PN) only adds a regularization item to the loss function in the training process of HRED, comparison with HRED model is sufficient for verifying the performance of the maximum entropy regularizer. From the results, HSCJN(w/o PN) achieves obvious improvement in both of the quality and diversity performance compared with HRED, demonstrating the effectiveness of the maximum entropy regularization. Both mechanisms contribute to the improvement of diversity and quality in response generation. The elimination of the prediction sub-network has greatest impact on the HSCJN model, indicating the importance of incorporating holistic semantic information.

4.6 Manual Evaluation

Since automatic metrics for open-domain generative models may not be consistent with human perceptions, the quality scores from the human annotations are more reliable. Therefore, we further recruit human annotators to evaluate the quality of the generated responses. We randomly select 100 testing dialogues with responses generated by different models for each dataset and for both 1-turn and 2-turns generation. Responses generated by different models are randomly shuffled for each annotator. 5 annotators with linguistics experience are recruited to refer to the test dialogue histories and judge the quality of the responses of all compared models according to the following criteria:

0: The response cannot be used as a response to the conversation context. It is semantically unrelated or disfluent.

+1: The response can be used as a reply to the message, but it is too universal like “Yes, I see.”, “Thank you.” and “I don’t know.”.

+2: The response is not only grammatical and relevant, but also informative and interesting.

Manual evaluation results for 1-turn response generation are presented in Table 3, which lists the percentage of each score and an overall average score. Among the three baselines, AttnSeq2Seq performs the worst and VHRED the best. Our model obtains a minimum of 0 points and a maximum of 2 points among all compared models on two datasets. This indicates that our model can generate fewer low-quality responses, as well as more semantically relevant and informative responses. The highest average score achieved by our model also confirms that our model outperforms the baselines.

Table 5: Examples of responses generated by different models given the multi-turn dialogue contexts

<p>Case 1: 1-turn response generation</p> <p>Speaker A: Good morning. Are you Mr.Liu? Speaker B: My name is Liu Lichi. How do you do? ...</p> <p>Speaker A: Have you had any working experience? Speaker B: Well, I worked at a supermarket during last summer holidays. Speaker A: How are your English and computer skills? Speaker B: I have passed the CET- 4 and 6. As far as computer is concerned, I can use the computer for word processing.</p>
<p>AttnSeq2Seq: A great idea. HRED: What do you do? VHRED: I think so. HSCJN: That sounds great. How long have you been interested in the job? Human: Okay. Mr.Liu, we'll inform you of the results within a week.</p>
<p>Case 2: 2-turns dialogue generation</p> <p>Speaker A: Can I borrow this magazine from you? It's really interesting and I can't put it down. Speaker B: I am sorry, but I can't lend it to you now, for I haven't finished reading it. If you don't mind, I can lend you some back numbers to you.</p>
<p>HRED: <i>Speaker A:</i> Thank you very much. <i>Speaker B:</i> You're welcome. VHRED: <i>Speaker A:</i> Thank you very much. <i>Speaker B:</i> Do you have any questions? HSCJN: <i>Speaker A:</i> That's great. <i>Speaker B:</i> It's too good as you like it. Human: <i>Speaker A:</i> That would be very kind of you. By the way, is it a monthly magazine? <i>Speaker B:</i> No, it is a fortnightly. So, you see, I can get the new one quite soon.</p>

Table 4 shows the manual evaluation results for 2-turns dialogue generation on DailyDialog dataset. In multi-turn generation, VHRED’s performance is not good. Responses generated by VHRED may be informative but most of them are irrelevant to the context. Our model also outperforms the baselines in terms of relevance and informativeness, as well as the overall average score.

5 Case Study

Table 5 presents the examples of 1-turn response and 2-turns dialogue generated by different models, given the multi-turn contexts between two speakers as inputs. “Human” lists the reference response in the dataset of the given context. We can see that Case 1 is an interview between an interviewer and a candidate. Our model captures that this is a job interview and produces a question matching the interview situation. In contrast, the responses generated by the baseline models are generic and irrelevant. In Case 2, our model captures the emotion information that speaker A likes the magazine, giving a more specific and informative response, and generates two consecutive turns matching different speakers’ roles, while the results generated by baselines are universal and monotone. It can be found that our model generates obviously better responses with more specific details and higher diversity. Moreover, it also shows that our results are more relevant to the dialogue scenario.

6 Conclusion

In this paper, we investigate the low diversity issue in dialogue generation task. We propose a Holistic Semantic Constraint Joint Network to introduce future information into the decoding stage. In addition, we devise a maximum entropy regularizer into our loss function to penalize the over-estimation of high frequency words. In this way, the model can see the entire target sequence and consider all the words in the vocabulary in the learning process. It is worth mentioning that our model introduces the language information in the target sequences from the Seq2Seq model itself for diverse response generation, which does not depend on any external information and variables, also beneficial to capture holistic dialogue semantics to promote relevance. Moreover, our joint learning framework can be generalized to any end-to-end model. Extensive experiments show that our model produces more informative and relevant responses than several competitive baselines.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Boxing Chen and Colin Cherry. 2014. A systematic comparison of smoothing techniques for sentence-level BLEU. In *WMT@ACL 2014, June 26-27, 2014, Baltimore, Maryland, USA*, pages 362–367.
- 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 *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5110–5117.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Gurvan L’Hostis, David Grangier, and Michael Auli. 2016. Vocabulary selection strategies for neural machine translation. *CoRR*, abs/1610.00072.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *NAACL HLT 2016, San Diego California, USA, June 12-17, 2016*, pages 110–119.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. A simple, fast diverse decoding algorithm for neural generation. *CoRR*, abs/1611.08562.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. In *IJCNLP 2017, Taipei, Taiwan, November 27 - December 1, 2017 - Volume 1: Long Papers*, pages 986–995.
- Kuan-Yen Lin, Chao-Chun Hsu, Yun-Nung Chen, and Lun-Wei Ku. 2019. Entropy-enhanced multimodal attention model for scene-aware dialogue generation. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence, (AAAI-19), Hilton Hawaiian Village, Honolulu, Hawaii, USA, January 27 C February 1, 2019*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1412–1421.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA.*, pages 311–318.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. 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, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 3776–3784.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. 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, February 4-9, 2017, San Francisco, California, USA.*, pages 3295–3301.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1577–1586.
- Yiping Song, Rui Yan, Yansong Feng, Yaoyuan Zhang, Dongyan Zhao, and Ming Zhang. 2018. Towards a neural conversation model with diversity net using determinantal point processes. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5932–5939.

- 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. In *NAACL HLT 2015, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 196–205.
- 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 *ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 2: Short Papers*, pages 231–236.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 2214–2218.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David J. Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *CoRR*, abs/1610.02424.
- Rongxiang Weng, Shujian Huang, Zaixiang Zheng, Xin-Yu Dai, and Jiajun Chen. 2017. Neural machine translation with word predictions. In *EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 136–145.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3351–3357.
- Jingjing Xu, Xu Sun, Xuancheng Ren, Junyang Lin, Bingzhen Wei, and Wei Li. 2018. DP-GAN: diversity-promoting generative adversarial network for generating informative and diversified text. *CoRR*, abs/1802.01345.
- Lili Yao, Yaoyuan Zhang, Yansong Feng, Dongyan Zhao, and Rui Yan. 2017. Towards implicit content-introducing for generative short-text conversation systems. In *EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2190–2199.
- Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. Augmenting end-to-end dialogue systems with commonsense knowledge. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4970–4977.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. 2018. Generating informative and diverse conversational responses via adversarial information maximization. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 1815–1825.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskénazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 654–664.