

Predict and Use Latent Patterns for Short-Text Conversation

Hung-Ting Chen*, Yu-Chieh Chao*, Ta-Hsuan Chao*, Wei-Yun Ma

Academia Sinica

{timchen0618, vpj870331, alexchao2007, ma}@iis.sinica.edu.tw

Abstract

Many neural network models nowadays have achieved promising performances in Chit-chat settings. The majority of them rely on an encoder for understanding the post and a decoder for generating the response. Without given assigned semantics, the models lack the fine-grained control over responses as the semantic mapping between posts and responses is hidden on the fly within the end-to-end manners. Some previous works utilize sampled latent words as a controllable semantic form to drive the generated response around the work, but few works attempt to use more complex semantic patterns to guide the generation. In this paper, we propose to use more detailed semantic forms, including latent responses and part-of-speech sequences sampled from the corresponding distributions, as the controllable semantics to guide the generation. Our results show that the richer semantics are not only able to provide informative and diverse responses, but also increase the overall performance of response quality, including fluency and coherence.

1 Introduction

The sequence-to-sequence neural network (Vinyals and Le 2015; Shang, Lu, and Li 2015) is one of the first successful neural models that generate text. However, the model are prone to producing generic and meaningless responses such as “I don’t know” or “So am I”.

To guide generation to produce more complex results, some works involve utilizing controllable semantics. Previous works (Ji, Lu, and Li 2014; Yan, Song, and Wu 2016; Yan et al. 2016) try to generate human-like and relevant responses by editing retrieved responses. However, the downside is that these models are not trained end-to-end, which have to retrieve candidate responses for every new input post in order to produce a response. On the other hand, numerous prior studies (Xing et al. 2016; Dziri et al. 2018; Mou et al. 2016; Yao et al. 2017; Gao et al. 2019a) use keywords, either selected by neural networks or retrieval models, as semantic features to control the generation of the output text. As these methods mainly use simple semantic forms to assist generation, the semantic restriction is relatively limited and lack the global view on a whole generated sentence.

*indicates equal contribution

In this paper, we propose an end-to-end model that creates informative and diverse responses by providing more complex, sentence-wide semantics. We investigate two different latent patterns - latent sentences and latent part-of-speech (POS) sequences for guiding semantics. First, as many previous works use a word to assist dialogue generation, we use a complete sentence instead. Sentences provide pattern and word usage information to yield robust dialogue responses. Second, inspired by Shen et al. (2019), who uses POS sequences to control the generation of Chinese poems, we serve POS sequences as hidden semantics to aid the Chit-chat task. We have shown that this method greatly helps improve the generation of dialogues, and that our generations indeed depend on the POS sequences we chose or generated.

Our model consists of two parts, the latent sequence predictor (including latent sentences and latent POS sequences), and the dialogue generator, which generates responses by the input post and the corresponding latent sequence we chose. After pre-training both parts, we combine the two networks into one end-to-end dialogue generation system. Finally, following Gao et al. (2019a), we fine-tune the whole network using a reinforcement learning (RL) algorithm. Our contributions can be summarized as follows:

- We propose an end-to-end model with RL that predicts and utilizes complex latent patterns to guide the generation of dialogues with given input text.
- We explore latent sentence and latent POS sequence as our semantic pattern to guide the generation, and find that the latter outperformed related state-of-the-art baselines.
- We implemented our dialogue generation models and all of our source codes and datasets are available at github.¹

2 Related Work

Many researchers have proposed methods to resolve the generic response problem. Maximum Mutual Information objective (Li et al. 2016a) and diverse beam search (Li, Monroe, and Jurafsky 2016) are used to address the issue. But these approaches only modify the generation at decoding steps. At training time, schemes such as reinforcement learning (Ranzato et al. 2015; Li et al. 2016b; Cuayáhuitl 2017) and adversarial training (Li et al. 2017; Xu et al. 2017) are

¹<https://github.com/timchen0618/LaPat>

also applied. One major drawback of these methods is that they cannot explicitly control what to produce .

Regarding control of semantics in a response, two type of methods prevail in the dialogue generation domain, as presented in the following subsections.

2.1 Rewriting and Editing a Sentence

Retrieval methods (Ji, Lu, and Li 2014; Yan, Song, and Wu 2016; Yan et al. 2016) have proved to produce more fluent outputs. Thus, many have proposed to rewrite the retrieved responses to balance between fluency and relevance to the input post. The candidates can serve as a good starting point, and then editing is performed to improve relevance.

For instance, Song et al. (2016) combine retrieval and generative systems by referencing a candidate response through multi-source attention and post-reranking. Weston, Dinan, and Miller (2018) refine the retrieved response by treating it as an additional context. Wu et al. (2019) treat the retrieved response as context and consider the lexical difference between the retrieved and the original post by augmenting the encoder-decoder model with an edit vector. Our work shares some similarities with these methods since we also reference a candidate response. However, we directly select a candidate response through neural models without doing retrieval first, eliminating the need for preprocessing of input data. Also, our model can be trained in an end-to-end manner and does inference as a whole.

2.2 Utilizing Latent Semantics Forms

Others have sought to enhance the quality of responses by utilizing latent features hidden in the post-response pairs since they provide content or functional information. A wide variety of methods are built upon the help of a latent word. Two previous studies (Xing et al. 2016; Dziri et al. 2018) select topic words with LDA (Latent Dirichlet Allocation) models and augment the encoder-decoder with topic-aware attention architecture. Mou et al. (2016) and Yao et al. (2017) both select a latent word with the highest PMI (Pointwise Mutual Information) against the input post and use it to guide generation. Gao et al. (2019a) first select a latent word from the vocabulary and make the decoder attend to both the latent word and the input post. Gao et al. (2019b) improve the CVAE structure by sampling a word from the vocabulary with a two-stage sampling scheme.

Researchers have also explored the possibility of using other latent semantic forms such as topic (Wang et al. 2017), sentence function (Ke et al. 2018; Bi et al. 2019), frame semantics (Gupta et al. 2020), and lexical phrases (Wu et al. 2020). Previous works have mainly focused on using much simpler semantic forms, thus altering few words in the responses and might not be as effective. Instead, we attempt to guide the generation with more complex semantic forms, explicitly providing the model with patterns to follow suit.

3 Models

The task we cope with in this paper is open-domain dialogue generation. An input post $\mathbf{p} \in \mathbf{P}$ is associated with multiple responses $\{\mathbf{r}\}$, where $\mathbf{r} \in \mathbf{R}$. Here \mathbf{P} and \mathbf{R} stands for the

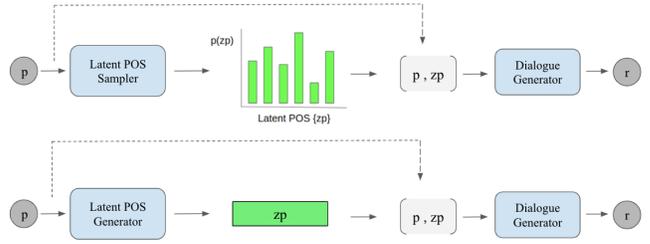


Figure 1: Overview of the proposed model. Latent sentence predictor share the architecture with latent POS sampler.

set of all the posts and the responses, respectively. We formulate the task as a sequence-to-sequence problem; given an input post $\mathbf{p} = \{p_1, \dots, p_i, \dots, p_n\}$ the model is expected to generate a response $\mathbf{r} = \{r_1, \dots, r_t, \dots, r_m\}$, where n and m represents the length of the post and response respectively. Latent semantic sequences $\mathbf{z} = \{z_1, \dots, z_j, \dots, z_l\}$ are introduced in the model. We denote latent sentences as $\mathbf{z}_s \in \mathbf{Z}_s = \mathbf{R}$ and latent part-of-speech tags as $\mathbf{z}_p \in \mathbf{Z}_p$, where $\mathbf{Z}_p = \{pos(\mathbf{r}) | \forall \mathbf{r} \in \mathbf{R}\}$.

Our model consists of two major components, as shown in figure 1. The latent sequence predictor aims to predict the latent semantic sequence \mathbf{z} conditioned on the post \mathbf{p} ; the dialogue generator produces a response \mathbf{r} according to \mathbf{p} and \mathbf{z} . Both models are first pretrained and then trained jointly in an end-to-end manner using reinforcement learning.

3.1 Latent Sequence Predictor

We tackle the selection of latent sentences and POS sequences as a classification problem. Due to the huge number of candidates, it is unlikely that we directly select one from the whole training dataset. Hence we construct a candidate set $\tilde{\mathbf{Z}}_i \subset \mathbf{Z}_i$ ($i \in \{s, p\}$) with K_i latent sequences from the latent space \mathbf{Z}_i , as described in section 3.3. Then for each post \mathbf{p} we sample a $\tilde{\mathbf{z}}_i$ from the candidate set $\tilde{\mathbf{Z}}_i$ to guide the generation. To produce more diverse POS sequences, we also propose a generation model for the Latent POS Predictor, which produces a POS sequence given the input post.

Latent Sentence Predictor This network samples latent sentence $\tilde{\mathbf{z}}_s$ from candidate set $\tilde{\mathbf{Z}}_s$ by estimating $p(\tilde{\mathbf{z}}_s | \mathbf{p})$. We first encode the input post \mathbf{p} with a bidirectional GRU (biGRU) to obtain an input representation \mathbf{h}^p , and retrieve hidden states of the last time step t of the biGRU encoder to represent the meaning of the post:

$$\mathbf{h}_t^p = [\overrightarrow{\mathbf{h}}_t^p, \overleftarrow{\mathbf{h}}_t^p] \tag{1}$$

and then compute the probability of each latent sentence:

$$p(\tilde{\mathbf{z}}_s | \mathbf{p}) = softmax(\text{MLP}(\mathbf{h}_t^p)) \tag{2}$$

where parameters in the biGRU layer and Multilayer Perceptron (MLP) are trainable. h_t^p denotes the hidden vector of last time step t of the encoder.

Latent POS Predictor Just like the latent sentence sampler, POS sequences can also be sampled from candidate set

$\tilde{\mathbf{Z}}_{\mathbf{p}}$. However, we also used a generation model for latent POS sequences to increase diversity. Next, we will elaborate on both the latent POS sampler and the latent POS generator.

- Latent POS Sampler

To sample a POS sequence, we simply use the Transformer (Vaswani et al. 2017) encoder to encode sentence:

$$\mathbf{h}^{\mathbf{P}} = \text{TransformerEncoder}(\mathbf{p}) \quad (3)$$

$$\mathbf{h}^{\mathbf{P}} = [\mathbf{h}_1^{\mathbf{P}}, \mathbf{h}_2^{\mathbf{P}}, \dots, \mathbf{h}_n^{\mathbf{P}}]$$

where n is the length of the input post. Then we treat the last time step of the hidden vectors $h_n^{\mathbf{P}}$ as our sentence representation. Finally, a MLP classifier is used to sample from the probability distribution of latent POS space:

$$p(\tilde{\mathbf{z}}_{\mathbf{p}}|\mathbf{p}) = \text{softmax}(\text{MLP}(\mathbf{h}_n^{\mathbf{P}})) \quad (4)$$

where parameters in the MLP block are trainable.

- Latent POS Generator

Instead of sampling a POS sequence from all POS sequence candidates, we adopt the Transformer to generate a POS sequence $\tilde{\mathbf{z}}_{\mathbf{p}\mathbf{g}}$ based on the input post \mathbf{p} .

3.2 Dialogue Generator

This model aims to generate response \mathbf{r} from predicted latent semantics and input post \mathbf{p} .

Dialogue Generator with Latent Sentence With a sampled $\tilde{\mathbf{z}}_s$ from latent sentence predictor, we first encode both the input post \mathbf{p} and the latent sentence $\tilde{\mathbf{z}}_s$ through two independent bidirectional GRU network for input representations $\mathbf{h}_{\mathbf{p}}^{\mathbf{g}}$ and $\mathbf{h}_{\tilde{\mathbf{z}}_s}^{\mathbf{g}}$ respectively. Both representations are then leveraged to decode the output sequence \mathbf{y} .

At each time step t , we add attention implemented as in (Bahdanau, Cho, and Bengio 2014):

$$e_i^t = \mathbf{v}^T \tanh(\mathbf{W}_h \mathbf{h}_i^{\mathbf{g}} + \mathbf{W}_s \mathbf{s}_t + b_{\text{attn}}) \quad (5)$$

$$\alpha_t = \text{softmax}(\mathbf{e}_t) \quad (6)$$

$$\mathbf{c}_t = \sum_i \alpha_i^t \mathbf{h}_i^{\mathbf{g}} \quad (7)$$

where \mathbf{v} , \mathbf{W}_h , \mathbf{W}_s and b_{attn} are learnable parameters, and \mathbf{c}_t is the context vector. The decoder attends to encoding representation of both the input post and the latent sentence. The above context vector calculation works for both sources.

We adopt pointer generator network (See, Liu, and Manning 2017). We calculate generation probability p_{gen} as:

$$p_{\text{gen}} = \sigma(\mathbf{W}_{\text{gen}}[\mathbf{s}_t, \mathbf{c}_t^{\mathbf{p}}, \mathbf{c}_t^{\tilde{\mathbf{z}}_s}]) \quad (8)$$

$$l_{\text{copy}} = 1 - p_{\text{gen}} \quad (9)$$

where l_{copy} denotes copy probability (copying word from the latent sentence). \mathbf{W}_{gen} is a learnable parameter, and σ is the sigmoid function. $\mathbf{c}_t^{\mathbf{p}}$ represents the input post context vector and $\mathbf{c}_t^{\tilde{\mathbf{z}}_s}$ represents the latent sentence context vector.

Finally, the probability distribution over the extended vocabulary (pre-set vocabulary & OOV words) is:

$$P(w) = p_{\text{gen}} P_{\text{vocab}}(w) + (1 - p_{\text{gen}}) \sum_{i:w_i=w} a_{i,t}^{\tilde{\mathbf{z}}_s} \quad (10)$$

where $a_{i,t}^{\tilde{\mathbf{z}}_s}$ denotes latent sentence attention distribution. Note that if w is an out-of-vocabulary (OOV) word, then $P_{\text{vocab}}(w)$, which represents the probability distribution over pre-set vocabulary, is zero.

Dialogue Generator with POS sequence Inspired by (Shen et al. 2019), we concatenate the predicted POS sequence $\tilde{\mathbf{z}}_{\mathbf{p}}$ (or in the generated POS sequence case, $\tilde{\mathbf{z}}_{\mathbf{p}\mathbf{g}}$) right behind the post at the input layer. Next, we forward the concatenated form into Transformer architecture as in (Vaswani et al. 2017) to generate the corresponding response:

$$\hat{\mathbf{r}} = \text{Transformer}([\mathbf{p}, \tilde{\mathbf{z}}_{\mathbf{p}}]) \quad (11)$$

The technique is capable of producing sentences matching the input POS sequences, hence providing enough guidance for generating responses with valid structure.

3.3 Pretraining and Data Preparation

We first pretrain the latent sentence predictor and the dialogue generator separately. For the latent sentence predictor and latent POS sampler, they are trained to solve a sentence classification problem. The creation of data for pretraining is explained in details below:

Candidate Set of Latent Sentences We attempt to construct a candidate set of size K_s representative enough of the entire response set \mathbf{R} . First, we use kmeans-clustering to aggregate responses into C clusters, and select K_s/C responses from each cluster. We then assign an index, i.e., label for classification, to each of the latent responses in $\tilde{\mathbf{Z}}_s$. Finally, we assign the label of $\tilde{\mathbf{z}}_s^*$ to each (\mathbf{p}, \mathbf{r}) , where $\tilde{\mathbf{z}}_s^*$ is the most similar response in $\tilde{\mathbf{Z}}_s$ with the target response \mathbf{r} . The similarity between responses is defined as the Euclidean distance of the BERT² sentence encoding of the two.

Candidate Set of Latent POS Sequences First, we derive the part-of-speech tagging of each response in \mathbf{R} using JIagu³. Likewise, the candidate POS sequences must be representative enough of $\mathbf{Z}_{\mathbf{p}}$. Thus we pick the K_p most common POS sequences in $\mathbf{Z}_{\mathbf{p}}$ as our candidate set $\tilde{\mathbf{Z}}_{\mathbf{p}}$.

Then, we create the training data for the latent POS sampler. We first assign classification label to each of the $\mathbf{z}_{\mathbf{p}} \in \tilde{\mathbf{Z}}_{\mathbf{p}}$. For each input post-response pair (\mathbf{p}, \mathbf{r}) , we find the most similar $\mathbf{z}_{\mathbf{p}}^*$ to the POS sequence of \mathbf{r} and label the pair with the corresponding index of $\mathbf{z}_{\mathbf{p}}^*$. That is, we wish to predict the class label of $\mathbf{z}_{\mathbf{p}}^*$ given the input post \mathbf{p} . The similarity between POS sequences is calculated by alignment.⁴

The generation of POS sequences and responses are both considered sequence-to-sequence generation problems. The

²<https://github.com/hanxiao/bert-as-service>

³<https://github.com/ownthink/JIagu>

⁴<https://biopython.org/DIST/docs/api/Bio.pairwise2-module.html>

latent semantic generator reads in the input post and outputs a POS sequence. For the dialogue generator, we simply concatenate the POS sequence \mathbf{z}_p of response \mathbf{r} after the post \mathbf{p} and input the newly created sequence $[\mathbf{p}; \mathbf{z}_p]$ into the Seq2Seq model.

3.4 Joint Training With Reinforcement Learning

To acquire better latent semantic sequences for the generation model, we fine-tune both models end-to-end with a reinforcement-learning objective. We treat the latent sequence predictor as an agent and the predictions of latent semantics as actions. We intend to choose the policy which enables the model to gain maximum total rewards. Note that during joint training, we apply the reinforcement learning algorithm only on the latent sequence predictor, whereas the dialogue generator is directly optimized through standard cross-entropy loss. We apply the REINFORCE (Williams 1992) algorithm, a Monte-Carlo based policy gradient method; the formula of updating parameters of the latent sequence predictor (denoted as θ) is as follows. $Q(\hat{\mathbf{z}})$ stands for estimated return if we sample a latent $\hat{\mathbf{z}}$ ($\hat{\mathbf{z}}$ can be either a latent sentence or a latent POS sequence), and we only calculate the return after the full response is generated.

- For selecting a sequence (either a sentence or a POS sequence)

$$\theta \leftarrow \theta + \nabla_{\theta} Q(\hat{\mathbf{z}}) \log p_{\theta}(\hat{\mathbf{z}}|\mathbf{p}) \quad (12)$$

- For generating a POS sequence

$$\theta \leftarrow \theta + \sum_{j=0}^l \nabla_{\theta} Q(\hat{\mathbf{z}}^j) \log p_{\theta}(\mathbf{z}_j^{\hat{\mathbf{p}}}|\mathbf{p}) \quad (13)$$

We model the estimated return as a reward function measuring the similarity between the predicted and target response. We update the latent sequence predictor only according to the maximum reward acquired from the bag of responses. We found empirically that it is challenging for the model to match all the responses during joint training since it has to explore the large latent space to select appropriate latent semantic sequences for all the responses. The reward function $Q(\hat{\mathbf{z}})$ is designed as follows; here $\{\mathbf{r}\}$ denote the set of multiple target responses associated with the post \mathbf{p}

$$Q(\hat{\mathbf{z}}) = R(\hat{\mathbf{r}}, \{\mathbf{r}\}) = \max_{\mathbf{r} \in \{\mathbf{r}\}} \mathbf{R}(\hat{\mathbf{r}}, \mathbf{r}) \quad (14)$$

, where

$$\mathbf{R}(\hat{\mathbf{r}}, \mathbf{r}) = F1(\hat{\mathbf{r}}, \mathbf{r}) \quad (15)$$

Here, a variation of F1 score calculates the overlap between the generated response and the target. Results have shown the joint training scheme enhances overall response quality.

4 Experiments

4.1 Dataset

Our method is evaluated on the Weibo Benchmark Dataset, an open-domain Chinese dialogue dataset with over 400 million training pairs. The testing set includes 3,200 data pairs. Here, we use the modified version released by (Gao et al. 2019a).⁵

⁵<https://ai.tencent.com/ailab/nlp/dialogue.html>

4.2 Evaluation

We report the results of the BLEU score (Papineni et al. 2002), a widely used metric measuring the similarity between the predicted response and the ground truth. We provide the values of BLEU 1-4 for all the experiments.

However, since the automatic evaluation metrics cannot effectively reflect the goodness of dialogue generation, we also did human evaluation on 100 randomly selected sentences from our and other comparing models. Three native speakers are asked to label the responses according to fluency, relevance and informativeness (generic responses would get low score in this category). Each response is given a score of 1 to 5 according to these quality factors.

4.3 Baseline Methods

We compare our method with popular models as baseline methods. The vanilla Seq2Seq model (Vinyals and Le 2015); HGFU (Yao et al. 2017), a model which modifies GRU by incorporating a cue word in the generation process; ResponseEdit (Wu et al. 2019), a retrieve-and-edit model utilizing an edit vector; CVAE (Zhao, Zhao, and Eskenazi 2017), and GMDR (Gao et al. 2019a), a latent word inference network combined with a joint-attention generation network.

4.4 Implementation Details

Data Preparation

- Latent Sentence

We set $C = 1000$ and $K_s = 50000$ for the construction of the candidate set $\tilde{\mathbf{Z}}_s$. The pretraining data for the latent sentence sampler thus includes input post \mathbf{P} as input and distribution of \mathbf{Z}_s^* as output.

- Latent POS Sequence

We expect the candidate set $\tilde{\mathbf{Z}}_p$ to be representative enough of \mathbf{Z}_p . In our experiments, we choose $K_p = 500, 1000, \text{ and } 10000$ to test how different size of candidate set affect the performance of our model.

Training Details

- Models Using Latent Sentences

In latent sentence models, we utilize GRU as our fundamental models. For the latent sentence predictor, we adopt 1-layer biGRU as our sentence encoder, and a classifier with 3 fully-connected layers. For the Dialogue Generator, we use 1-layer biGRU for the encoders of both input posts and latent sentences. A decoder of 1 layer uni-directional GRU is used. The dimensions are 1024 for the encoders and 512 for the decoder and classifier.

All models are trained with Adam optimizer. The learning rate for the latent sentence predictor and the dialogue generator are set to 0.002 and 0.0002, respectively. The learning rate decay is 0.5 each epoch. During testing, we apply beam-search with beam size 4.

- Models Using Latent POS Sequences

For experiment utilizing latent POS sequences, the dialogue generator is a standard 6-layer Transformer encoder-decoder model. We follow the model architecture and the parameter

Models	BLEU-1	BLEU-2	BLEU-3	BLEU-4
Seq2Seq	25.57	7.39	1.71	0.64
HGFU (Yao et al. 2017)	27.31	9.59	4.08	2.14
ResponseEdit (Wu et al. 2019)	24.45	7.32	2.38	0.96
CVAE (Zhao, Zhao, and Eskenazi 2017)	20.13	6.82	1.36	0.52
GMDR (Gao et al. 2019a)	21.26	8.27	3.93	2.08
Ours-Two Stage (latent sentence)	17.19	4.72	0.42	0.10
Ours-Two Stage (sample POS)	22.28	5.31	1.11	0.00
Ours-Two Stage (generate POS)	22.84	6.08	1.07	0.00
Ours-RL (latent sentence)	19.21	6.83	1.61	0.23
Ours-RL (sample POS)	36.42	12.35	3.55	0.99
Ours. RL (generate POS)	30.98	10.21	3.25	1.29

Table 1: Automatic evaluation results on Weibo.

of the base model in Vaswani et al. (2017). The latent POS sampler model is a 6-layer transformer encoder with a classifier the same as the Latent Sentence Predictor on top of the last time step of hidden states. The latent POS generator shares the model architecture with the dialogue generator.

All Transformer models are trained with the Adam optimizer. For the pretrained models, we adopt noam learning rate decay as proposed by Vaswani et al. (2017); the warmup step is set to 8000. The initial learning rate is set to 10^{-5} for joint fine-tuning, and is decayed every epoch. During testing, the beam size is set to 3.

5 Results

5.1 Automatic Evaluation

The evaluation results of our model are shown at Table 1. As we can see, our Ours-RL(sample POS) outperforms other models in BLEU 1 and 2. The reason is that the sentence patterns we provided narrow down the word selection at each time step. For example, if the corresponding POS tag for the token at this time step is v , then the model will generate a verb with higher probability. Thus, our model can choose more accurate and relevant words, and thus achieves outstanding performance in BLEU 1 and 2. However, since our model forces the pattern of our generation (which is normally different from the original response), it does not perform as well in longer word matching (BLEU 3 and 4).

We also report the performance of combining the pretrained models as Ours-Two Stage, i.e., directly using the latent sequence predicted by pretrained latent sequence predictor during generation. The BLEU scores improved significantly after fine-tuning with our reinforcement learning algorithm since both modules are optimized in an end-to-end manner, enabling more appropriate prediction of latent sequences. The BLEU score of Ours-RL(sample POS) improved **63%** and **132%** and achieved **36.42** and **12.35** in BLEU-1 and BLEU-2 scores.

However, our latent sentence model does not yield good BLEU scores because the generation is greatly influenced by the chosen latent sentence, as shown in Table 5. Thus, if the latent sentence we have chosen is very different from the original response, the BLEU scores would drop. Also, the search space for the latent sentence sampler is way smaller

Models	Fluency	Relevance	Informativeness
HGFU	3.69	2.65	2.60
GMDR	4.51	3.03	2.73
Ours(select POS)	3.98	3.36	3.34

Table 2: Result of Human Evaluation.

than the ResponseEdit model, which uses pre-define index to search for similar posts in all posts at test time. Possible future work could be increasing the latent space and add some pre-defined constraints during decoding.

5.2 Human Evaluation

The result of human evaluation is shown in Table 2. The model that we use for comparison is latent POS sampler and $K_p = 500$. Our model obtains the highest scores in both relevance and informativeness. The reason is that our model incorporates latent patterns, which provide guidance signal about how to organize content and when to include certain information. With a reasonable pattern to follow, the generated responses are thus more relevant and informative. However, since the outputs of the generated sentences follow the POS patterns, they might sacrifice fluency slightly. The fluency score is still better than HGFU, though.

5.3 Case Study and Analysis

Table 3 shows one of the dialogue predicted by our models and other models. The model GMDR (Gao et al. 2019a) predicted a latent word *darling*, and thus generated a simple sentence *All of them are my darlings*. However, our models produced more complex results, both in word usage and sentence patterns. Our latent sentence model chose *The sky of the liberated area is a sunny day*. as the latent sentence. (The sentences in brackets are selected latent responses.) Thus, it affected the generation *Oh my god* (because in Chinese, *day* 天 is equal to *god* 天). As for the latent-POS-predicting models, the selected POS sequence provided a reasonable output sentence pattern for the model to follow. Thus, we obtain an informative and relevant response *They are so cute, I'll be blessed, too. I miss my darling so much*.

Input	这些都是我的宝贝们 All of these are my darlings.
Seq2Seq	我也想说, 我也想吃。 I also want to say; I want to eat, too.
HGFU	都是我喜欢的, 好可爱 All of them are my favorite, they're so cute.
GMDR	都是我的宝贝们。 All of them are my darlings.
CVAE	还有那个可以吗? Is that available?
Ours(Latent Sentence)	我的天我的天, 喜欢。 [解放区的天是晴朗的天] Oh my god oh my god, like it. [The sky of the liberated area is a sunny day.]
Ours(Generate POS)	好可爱啊, 好可爱, 好喜欢。 So cute, so cute, and I like it so much.
Ours(Sample POS)	好可爱啊, 我也要幸福的。好想我的宝宝啊。 They are so cute, I'll be blessed, too. I miss my darling so much.

Table 3: Evaluation results of different sizes of latent space

K_p	BLEU-1	BLEU-2	BLEU-3	BLEU-4
500	36.42	12.35	3.55	1.29
1000	28.60	9.12	3.09	1.07
10000	22.17	7.51	2.70	1.03

Table 4: Evaluation results of different sizes of latent space

5.4 Analysis On the Selection of Latent Space

Table 4 shows the results of how different latent space size K_p affect the performance of our model. We can see that the best results occur at $K_p = 500$. If the latent spaces is too large, then it is difficult for the model to search through the whole space for appropriate latent sequences.

5.5 Faithfulness to Given Semantic Sequences

To prove that the selected latent semantic sequence is actually helpful for generation, we inspect if the generated response follows the provided word or POS sequence.

For model utilizing latent sentences, we inquire whether the model mimics the word selection (providing content) and the word ordering (providing word usage, sentence pattern, and content ordering) of the given latent sentence. As for word selection, we calculate the percentage of overlapping words with the latent sentences in generated responses. For word order, we calculate the percentage of overlapping n-grams ($n \geq 2$). The result is presented in table 5.

As can be seen from the result, the model copies over two-third of words and about one-fourth of the words from the latent sentence. It clearly illustrates that the model indeed ref-

	unigram	bigram	trigram	4-gram
overlap (%)	71.11	44.86	32.84	26.76

Table 5: N-gram overlap - latent sentences and responses

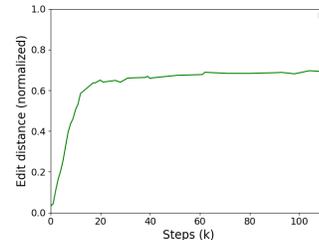


Figure 2: Edit distance during fine-tuning process of RL.

erence to the latent sentence for better response generation. However, the model might also suffers from copying bad latent sentences. Thus we believe the overall response quality will improve if we could ensure better latent sentence selection, which may be a potential future study.

For models using latent POS sequences, the output responses mostly follow the given POS sequence during pre-training but learns to generate more freely when fine-tuning. We believe our model not only follows the guidance of the underlying patterns of latent POS, but also maintains a certain extent of freedom concerning the generation of word sequences. To support the claim, we measure the faithfulness of generation to the provided POS sequence using the edit distance between the POS of output responses and the given POS sequence. Smaller edit distance indicates greater faithfulness to the latent POS. The edit distance is normalized by the length of the selected POS sequence to simulate the mismatch percentage.

Figure 2 shows the result during fine-tuning for our best model. The model first strictly follows the selected latent POS sequence but later begins to generate responses more freely. We thus prove statistically that our model indeed strikes a balance between generating word segments on its own and referencing to the latent POS sequence. The automatic evaluation results also proves our theory, since fine-tuned models greatly outperform the two-stage models.

6 Conclusion

We propose an end-to-end response generation network aided by complex latent semantics. Our model uses latent sentences and latent POS sequences respectively to provide wording and sentence pattern information. Also, our combined end-to-end model can be optimized by a reinforcement learning algorithm, and is capable of striking a balance between totally and partially following the latent semantics during generation. Our results show that our model increases the quality of responses by being more informative and relevant to the input post compared with existing baselines.

7 Future Work

Our next step would be using a mixed latent sequence to help guide the generation. For example, we can integrate both POS tags and words into a sequence to help our model generate more informative and human-like results. Also, we could try other semantic forms to assist our generation, such as Semantic Role Labeling (SRL).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Bi, W.; Gao, J.; Liu, X.; and Shi, S. 2019. Fine-grained sentence functions for short-text conversation. *arXiv preprint arXiv:1907.10302*.
- Cuayáhuitl, H. 2017. SimpledS: A simple deep reinforcement learning dialogue system. In *Dialogues with social robots*, 109–118. Springer.
- Dziri, N.; Kamaloo, E.; Mathewson, K. W.; and Zaiane, O. 2018. Augmenting neural response generation with context-aware topical attention. *arXiv preprint arXiv:1811.01063*.
- Gao, J.; Bi, W.; Liu, X.; Li, J.; and Shi, S. 2019a. Generating multiple diverse responses for short-text conversation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 6383–6390.
- Gao, J.; Bi, W.; Liu, X.; Li, J.; Zhou, G.; and Shi, S. 2019b. A discrete cvae for response generation on short-text conversation. *arXiv preprint arXiv:1911.09845*.
- Gupta, P.; Bigham, J. P.; Tsvetkov, Y.; and Pavel, A. 2020. Controlling Dialogue Generation with Semantic Exemplars. *arXiv preprint arXiv:2008.09075*.
- Ji, Z.; Lu, Z.; and Li, H. 2014. An information retrieval approach to short text conversation. *arXiv preprint arXiv:1408.6988*.
- Ke, P.; Guan, J.; Huang, M.; and Zhu, X. 2018. Generating informative responses with controlled sentence function. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1499–1508.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, B. 2016a. A Diversity-Promoting Objective Function for Neural Conversation Models. In *Proceedings of NAACL-HLT*, 110–119.
- Li, J.; Monroe, W.; and Jurafsky, D. 2016. A simple, fast diverse decoding algorithm for neural generation. *arXiv preprint arXiv:1611.08562*.
- Li, J.; Monroe, W.; Ritter, A.; Galley, M.; Gao, J.; and Jurafsky, D. 2016b. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541*.
- Li, J.; Monroe, W.; Shi, T.; Jean, S.; Ritter, A.; and Jurafsky, D. 2017. Adversarial learning for neural dialogue generation. *arXiv preprint arXiv:1701.06547*.
- Mou, L.; Song, Y.; Yan, R.; Li, G.; Zhang, L.; and Jin, Z. 2016. Sequence to backward and forward sequences: A content-introducing approach to generative short-text conversation. *arXiv preprint arXiv:1607.00970*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Ranzato, M.; Chopra, S.; Auli, M.; and Zaremba, W. 2015. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. *arXiv preprint arXiv:1503.02364*.
- Shen, L.-H.; Tai, P.-L.; Wu, C.-C.; and Lin, S.-D. 2019. Controlling Sequence-to-Sequence Models-A Demonstration on Neural-based Acoustic Generator. 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): System Demonstrations*, 43–48.
- Song, Y.; Yan, R.; Li, X.; Zhao, D.; and Zhang, M. 2016. Two are better than one: An ensemble of retrieval-and generation-based dialog systems. *arXiv preprint arXiv:1610.07149*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Vinyals, O.; and Le, Q. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Wang, D.; Jovic, N.; Brockett, C.; and Nyberg, E. 2017. Steering output style and topic in neural response generation. *arXiv preprint arXiv:1709.03010*.
- Weston, J.; Dinan, E.; and Miller, A. H. 2018. Retrieve and refine: Improved sequence generation models for dialogue. *arXiv preprint arXiv:1808.04776*.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4): 229–256.
- Wu, Y.; Wei, F.; Huang, S.; Wang, Y.; Li, Z.; and Zhou, M. 2019. Response generation by context-aware prototype editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7281–7288.
- Wu, Z.; Galley, M.; Brockett, C.; Zhang, Y.; Gao, X.; Quirk, C.; Koncel-Kedziorski, R.; Gao, J.; Hajishirzi, H.; Ostendorf, M.; et al. 2020. A controllable model of grounded response generation. *arXiv preprint arXiv:2005.00613*.
- Xing, C.; Wu, W.; Wu, Y.; Liu, J.; Huang, Y.; Zhou, M.; and Ma, W.-Y. 2016. Topic aware neural response generation. *arXiv preprint arXiv:1606.08340*.
- Xu, Z.; Liu, B.; Wang, B.; Sun, C.-J.; Wang, X.; Wang, Z.; and Qi, C. 2017. Neural response generation via gan with an approximate embedding layer. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 617–626.
- Yan, R.; Song, Y.; and Wu, H. 2016. Learning to respond with deep neural networks for retrieval-based human-computer conversation system. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, 55–64.
- Yan, R.; Song, Y.; Zhou, X.; and Wu, H. 2016. "Shall I Be Your Chat Companion?" Towards an Online Human-Computer Conversation System. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 649–658.
- Yao, L.; Zhang, Y.; Feng, Y.; Zhao, D.; and Yan, R. 2017. Towards implicit content-introducing for generative short-text conversation systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2190–2199.
- Zhao, T.; Zhao, R.; and Eskenazi, M. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. *arXiv preprint arXiv:1703.10960*.