

Adaptive Parameterization for Neural Dialogue Generation

Hengyi Cai^{†,§,*}, Hongshen Chen[†], Cheng Zhang[†], Yonghao Song[†], Xiaofang Zhao[†] and Dawei Yin[‡]

[†]Institute of Computing Technology, Chinese Academy of Sciences

[§]University of Chinese Academy of Sciences

[‡]Data Science Lab, JD.com

caihengyi@ict.ac.cn, ac@chenhongshen.com,

{zhangcheng, songyonghao, zhaoxf}@ict.ac.cn, yindawei@acm.org

Abstract

Neural conversation systems generate responses based on the sequence-to-sequence (SEQ2SEQ) paradigm. Typically, the model is equipped with a single set of learned parameters to generate responses for given input contexts. When confronting diverse conversations, its adaptability is rather limited and the model is hence prone to generate generic responses. In this work, we propose an **Adaptive Neural Dialogue** generation model, ADAND, which manages various conversations with conversation-specific parameterization. For each conversation, the model generates parameters of the encoder-decoder by referring to the input context. In particular, we propose two adaptive parameterization mechanisms: a context-aware and a topic-aware parameterization mechanism. The context-aware parameterization directly generates the parameters by capturing local semantics of the given context. The topic-aware parameterization enables parameter sharing among conversations with similar topics by first inferring the latent topics of the given context and then generating the parameters with respect to the distributional topics. Extensive experiments conducted on a large-scale real-world conversational dataset show that our model achieves superior performance in terms of both quantitative metrics and human evaluations.

1 Introduction

Open-domain dialogue models, usually called chat systems, draw increasing attention from both academia and industry in recent years. Building on the successful sequence-to-sequence (SEQ2SEQ) paradigm (Sutskever et al., 2014; Cho et al., 2014; Bahdanau et al., 2015), contemporary mainstream open-domain dialogue generation models (Shang et al., 2015; Serban et al., 2016, 2017; Shen et al., 2017; Clark and Cao, 2017; Zhao et al., 2017; Xing

et al., 2017; Zhao et al., 2018), trained on a large number of context-response pairs, attempt to generate an appropriate response for the given context based on a single set of the model parameters.

Because of its great potential in understanding and modeling conversations, SEQ2SEQ has been widely applied in different kinds of conversation scenarios including technical supporting, movie discussions, and social entertainment, etc. However, when confronting conversations with diverse topics or themes, SEQ2SEQ is usually prone to make generic meaningless responses due to its oversimplified parameterization. To tackle this issue, Xing et al. (2017) proposed a topic-aware response generation model, which utilizes a single encoder-decoder, augmented with topic information obtained from a pre-trained LDA model. Though effective, the model heavily relies on the outsourcing topic information to capture the topic variations of different conversations. Another approach, per-topic/theme encoder-decoder model (Choudhary et al., 2017), uses separate encoder-decoder model for each topic or theme. This method needs the preorganized topic/theme annotations for each conversation, which are prohibitively expensive to obtain. Furthermore, building multiple separate topic/theme-specific encoder-decoders not only weakens the applicability and efficiency of the system, but also prevents parameter sharing across domains, which leads to overparameterization due to the excessive amount of model parameters.

To gather the benefits of both approaches, in this paper, we propose an adaptive neural dialogue generation model which utilizes a single encoder-decoder for diverse conversations, meanwhile, the encoder-decoder is specifically parameterized according to each conversation. In particular, we propose two adaptive parameterization mechanisms: 1) context-aware parameterization directly generates parameters of the encoder-decoder model by capturing local semantics of the given context; 2)

*Work done at Data Science Lab, JD.com.

topic-aware parameterization enables parameter sharing among conversations with similar topic distributions by first inferring the latent topics of the input context, and then generating the parameters with respect to the inferred latent topics. Equipped with both the context-aware and topic-aware parameterization mechanisms, the model is capable of generating responses for diverse conversations with a single encoder-decoder through a more flexible and efficient approach. Moreover, our model is trained in an end-to-end fashion without any costly external or labeled topic annotations.

We empirically evaluate our approach on a large scale real-world conversational dataset. Extensive experiments show that our proposed ADAND outperforms existing dialogue generation models in terms of both the automatic evaluation metrics and human judgements.

2 Neural Dialogue Generation Model

Conventionally, neural dialogue generation model follows the encoder-decoder paradigm. Given the context $\mathbf{x} = \{x_1, x_2, \dots, x_{T_x}\}$ and the target response $\mathbf{y} = \{y_1, y_2, \dots, y_{T_y}\}$, the model learns to maximize the following conditional probability:

$$p(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^{T_y} p(y_t | \mathbf{y}_{<t}, \mathbf{x}), \quad (1)$$

where $\mathbf{y}_{<t} = y_1 \dots y_{t-1}$.

Typically, the probability $p(y_t | \mathbf{y}_{<t}, \mathbf{x})$ is modeled as follows:

$$p(y_t | \mathbf{y}_{<t}, \mathbf{x}) = p(y_t | s_{t-1}), \quad (2)$$

where s_{t-1} is the decoding hidden state up to time step $t - 1$, depending on $\mathbf{y}_{<t}$ and \mathbf{x} . The s_t can be defined as:

$$s_t = \mathcal{F}(y_t, s_{t-1}, \mathcal{G}(\mathbf{x})), \quad (3)$$

where the encoder \mathcal{G} and the decoder \mathcal{F} can be implemented as recurrent neural networks such as LSTMs (Hochreiter and Schmidhuber, 1997) or GRUs (Cho et al., 2014), or the transformer (Vaswani et al., 2017) (with the attention mechanism (Bahdanau et al., 2015) or self-attention mechanism).

In this work, we employ the LSTM-based encoder-decoder dialogue generation model. The




Figure 1: General model architecture. Black solid lines denote information flow, and the red dashed line indicates the adaptive parameterization operation.

LSTM unit is formulated as:

$$\begin{aligned} i_t &= W_h^i h_{t-1} + W_x^i x_t + b^i \\ g_t &= W_h^g h_{t-1} + W_x^g x_t + b^g \\ f_t &= W_h^f h_{t-1} + W_x^f x_t + b^f \\ o_t &= W_h^o h_{t-1} + W_x^o x_t + b^o \\ c_t &= \sigma(f_t) \odot c_{t-1} + \sigma(i_t) \odot \tanh(g_t) \\ h_t &= \sigma(o_t) \odot \tanh(c_t), \end{aligned} \quad (4)$$

where σ is the sigmoid operator and \odot stands for Hadamard product. h_{t-1} is the previous hidden state and x_t is the input embedding at step t . $W = \{W_h^i, W_h^g, W_h^f, W_h^o, W_x^i, W_x^g, W_x^f, W_x^o\}$ and $b = \{b^i, b^g, b^f, b^o\}$ are the LSTM parameters. The model parameters are tuned on the training corpus.

When testing, given the input context, it generates a response with the learned parameters. This architecture shows great success in neural dialogue generation (Shang et al., 2015; Serban et al., 2016, 2017; Shen et al., 2017; Clark and Cao, 2017; Xing et al., 2017; Choudhary et al., 2017; Zhao et al., 2018). However, with a single set of model parameters and the oversimplified model architecture, the flexibility of the model is rather limited, especially when confronting conversations with diverse topics or themes.

3 Adaptive Neural Dialogue Generation Model

In this work, we propose an adaptive neural dialogue generation model which utilizes a single encoder-decoder model and a set of dynamical parameters to balance the model's flexibility and efficiency. As depicted in Figure 1, we utilize the model adapter to parameterize the encoder-decoder for each conversation. It takes the dialogue context as its input, and generates parameters of the encoder-decoder model through two adaptive parameterization mechanisms; and then the resultant




Figure 2: General architecture of the context-aware parameterization. (a) Context-aware parameter adapter. (b) One layer of the dialogue generation model. Black solid lines denote information flow, and red dashed lines indicate adaptive parameterization operations. \mathcal{A}_c stands for the context-aware parameterization function (Eq.(7)).

encoder-decoder model generates the response with a specific set of model parameters.

3.1 Context-aware Parameterization

Context-aware parameterization parameterizes the dialogue encoder-decoder with respect to the local semantics of the given context. As shown in Figure 2, we first obtain the semantic representation of the context through a context encoder. Then, the parameter adapter dynamically adapts the parameters of the encoder-decoder at each time step. Here, we utilize a bidirectional LSTM to transform the input context into the semantic hidden representation ζ .

The parameter adapter then generates the weights of LSTM units as:

$$W = M^{[1:N_r]}\zeta, \quad (5)$$

where $W \in \mathbb{R}^{N_r \times N_c}$ and $M^{[1:N_r]} \in \mathbb{R}^{N_r \times N_c \times N_c}$ is a tensor. $N_r = 4N_h$ and $N_c = N_h + N_x$, in which N_h is the hidden size of the LSTM and N_x is the embedding size. $\zeta \in \mathbb{R}^{N_c}$ is the context representation. The product $M^{[1:N_r]}\zeta$ results in a weight matrix W , where each row is computed by one slice $j = 1, 2, \dots, N_r$ of the tensor: $W_j = (M^j \zeta)^T$.

Although such parameterization seems to be straightforward, due to the quadratic size of $M^{[1:N_r]}$, the parameter size of such parameter adapter is N_c times larger than the basic encoder-decoder model. Therefore, overfitting

and expensive computational cost make it infeasible (Bertinetto et al., 2016; Ha et al., 2017). Following Flennerhag et al. (2018), we reduce the parameter space by factorizing the weights as:

$$\begin{aligned} W &= \mathcal{A}_c(\xi_t, U, V) \\ \xi_t &= \Phi(\zeta, h_{t-1}), \end{aligned} \quad (6)$$

where Φ is implemented as a LSTM unit, h_{t-1} is the previous encoder or decoder hidden state, and $\xi_t \in \mathbb{R}^{N_c}$ is the context representation at time step t . \mathcal{A}_c denotes the context-aware parameterization function, defined as:

$$\mathcal{A}_c(\xi_t, U, V) = U\xi_t V^T, \quad (7)$$

where $U \in \mathbb{R}^{N_r \times N_c}$ and $V \in \mathbb{R}^{N_c \times N_c}$ are learnable weights.

The context-aware parameterization function \mathcal{A}_c is reminiscent of the Singular Value Decomposition (SVD). Here, \mathcal{A}_c composes a projection by adapting the dialogue context to the weight matrices and does not perform matrix decomposition actually. The number of parameters in this formulation is $L = N_r \times N_c + N_c \times N_c$ and the total parameter number of the model is linear with L , so that the total number of model parameters does not explode.

3.2 Topic-aware Parameterization

Context-aware parameterization adapts the encoder-decoder parameters regarding each input context. As a result, the adapted encoder-decoder is sensitive to the input context. To enable the parameter sharing among similar topics, we further propose a topic-aware parameterization mechanism. As shown in the Figure 3, the topic inferrer first distills the topic distribution θ from the context (Figure 3 (a)), and then the parameters of the encoder-decoder model are constructed upon θ (Figure 3 (b)).

3.2.1 Latent Topic Inference

We introduce a variational topic inferrer to infer the topic distribution θ of the conversation. Drawing inspirations from neural topic model (Miao et al., 2017), as illustrated in Figure 3 (a), the generation process of the variational topic inferrer is formulated as follows:

- (i) A latent variable ν is inferred to convey the underlying semantics of the given context.




Figure 3: Illustration of the topic-aware parameterization. (a) Topic inferrer. (b) One layer of the dialogue generation model. Black solid lines denote information flow, blue dashed lines only appear in training, and red dashed lines indicate adaptive parameterization operations.

- (ii) The topic distribution θ is constructed from the latent variable ν through a softmax function.
- (iii) The dialogue \mathbf{d} , composing of a context-response pair (\mathbf{x}, \mathbf{y}) , is drawn from the topic distribution over words $p(w_i|\beta_{z_i})$, where z_i is a topic assignment sampled from a multinomial distribution parameterized by the topic distribution θ , and β_{z_i} is the topic-word distribution of topic assignment z_i .

Given a context \mathbf{x} , the latent variable ν is sampled from $P(\nu|\mathbf{x}) = \mathcal{N}(\mu_{prior}, \sigma_{prior}^2)$, and $\mathcal{N}(\mu_{prior}, \sigma_{prior}^2)$ is the multivariate normal distribution with mean μ_{prior} and diagonal covariance matrix σ_{prior}^2 . In practice, ν is reparameterized as: $\nu = \mu_{prior} + \epsilon \cdot \sigma_{prior}$ and ϵ is a standard Gaussian noise. The Gaussian parameters μ_{prior} and σ_{prior}^2 are the outputs of multilayer perceptrons (MLP) given the bag-of-words (BoW) representation of the context as input. To reduce the encoding noise of stop words, here we choose the BoW instead of LSTM for context representations, following Miao et al. (2017). To implement neural variational inference, we utilize an inference network $Q(\nu|\mathbf{d}) = \mathcal{N}(\mu_{posterior}, \sigma_{posterior}^2)$ to approximate the intractable true posterior $p(\nu|\mathbf{d})$,

where $\mu_{posterior}$ and $\sigma_{posterior}^2$ are computed in the same way as the prior, taking the bag-of-words representation of dialogue \mathbf{d} as input. The dialogue \mathbf{d} consists of the context-response pair (\mathbf{x}, \mathbf{y}) .

The topic distribution θ is constructed from the latent variable ν through a softmax function:

$$\theta = g(\nu) = \text{softmax}(\nu \cdot \mathbf{W}_\nu), \quad (8)$$

where \mathbf{W}_ν is a linear transformation and the bias terms are left out for brevity.

Then, the dialogue \mathbf{d} is generated based on the topic distribution θ . Given θ , the marginal likelihood of a dialogue \mathbf{d} is formulated as:

$$p(\mathbf{d}) = \int_{\theta} p(\theta) \prod_{i=1}^{|\mathbf{d}|} \sum_{z_i} p(w_i|\beta_{z_i}) p(z_i|\theta) d\theta. \quad (9)$$

In addition, the topic assignment z_i can be integrated out and the log-likelihood of a word w_i in dialogue \mathbf{d} can be factorized as:

$$\begin{aligned} \log p(w_i|\beta, \theta) &= \log \sum_{z_i} [p(w_i|\beta_{z_i}) p(z_i|\theta)] \\ &= \log(\theta \cdot \beta^T). \end{aligned} \quad (10)$$

The topic-word distribution β_k is defined by:

$$\beta_k = \text{softmax}(F \cdot \Lambda_k^T), \quad (11)$$

where $F \in \mathbb{R}^{C \times H}$ is the topical word embedding matrix and $\Lambda \in \mathbb{R}^{K \times H}$ is the topic embedding matrix, K is the number of topics, C is the number of topical words, H is the embedding size and $\beta = \{\beta_1, \beta_2, \dots, \beta_K\}$.

3.2.2 Parameterization with Latent Topics

We parameterize the encoder-decoder with the inferred topic distribution θ . In context-aware parameterization, the parameters of the encoder-decoder are adapted dynamically at each time step, whereas in topic-aware parameterization, as illustrated in Figure 3 (b), we generate only one set of parameters for each conversation.

Similar to the context-aware parameterization function in Eq.(7), given the topic distribution θ , the topic-aware parameterization function \mathcal{A}_κ constructs the LSTM weight W as follows:

$$\begin{aligned} W &= \mathcal{A}_\kappa(\theta, U, V) \\ \mathcal{A}_\kappa(\theta, U, V) &= U\theta V^T, \end{aligned} \quad (12)$$

where $U \in \mathbb{R}^{N_r \times K}$ and $V \in \mathbb{R}^{N_c \times K}$ are learnable parameters. K is the number of latent topics.

3.3 Parameterization with Both Context and Topics

Intuitively, context-aware parameterization is more adept at capturing local semantics of the input context while topic-aware parameterization enables parameter sharing between conversations with similar topic distributions. To benefit the model parameterization with both the local and global information, we further adapt parameters of the encoder-decoder by utilizing both the context representations ξ_t and the topic distribution θ . In particular, the LSTM weight W at time step t is adapted as follows:

$$\begin{aligned} W &= \Psi_t \mathcal{A}_c(\xi_t, U_c, V_c) + \\ &\quad (1 - \Psi_t) \mathcal{A}_\kappa(\theta, U_\kappa, V_\kappa) \quad (13) \\ \Psi_t &= \sigma(\xi_t, \theta), \end{aligned}$$

where Ψ_t is the gating function deciding whether the parameterization relies more on the context or the topics. U_c, V_c, U_κ and V_κ are learnable weights. \mathcal{A}_c and \mathcal{A}_κ denote the context-aware and topic-aware parameterization function respectively. σ is the sigmoid function.

3.4 Learning

To enable the joint optimization of latent topic inference, adaptive model parameterization, and response generation in ADAND, given the definitions in Eq.(9), similar to Kingma and Welling (2014) and Miao et al. (2017), we derive a variational lower bound for the generation likelihood:

$$\begin{aligned} \mathcal{J} &= \mathbb{E}_{Q(\theta|\mathbf{d})} \left[\sum_{i=1}^{|\mathbf{d}|} \log \sum_{z_i} [p(w_i | \beta_{z_i}) p(z_i | \theta)] \right] + \log p(\mathbf{y}|\mathbf{x}) \\ &\quad - D_{KL}(Q(\theta|\mathbf{d}) || P(\theta|\mathbf{x})) \\ &= \mathbb{E}_{Q(\theta|\mathbf{d})} \left[\sum_{i=1}^{|\mathbf{d}|} \log p(w_i | \beta, \theta) + \sum_{t=1}^{T_y} \log p(y_t | s_{t-1}) \right] \\ &\quad - D_{KL}(Q(\theta|\mathbf{d}) || P(\theta|\mathbf{x})) \\ &= \mathbb{E}_{Q(\nu|\mathbf{d})} \left[\sum_{i=1}^{|\mathbf{d}|} \log p(w_i | \beta, \nu) + \sum_{t=1}^{T_y} \log p(y_t | s_{t-1}) \right] \\ &\quad - D_{KL}(Q(\nu|\mathbf{d}) || P(\nu|\mathbf{x})) \\ &\approx \sum_{i=1}^{|\mathbf{d}|} \log p(w_i | \beta, \nu) + \sum_{t=1}^{T_y} \log p(y_t | s_{t-1}) \\ &\quad - D_{KL}(Q(\nu|\mathbf{d}) || P(\nu|\mathbf{x})), \end{aligned} \quad (14)$$

where $\mathbf{y} = \{y_1, y_2, \dots, y_{T_y}\}$, $P(\nu|\mathbf{x})$ is the prior estimation of the latent variable ν which approximates the posterior $Q(\nu|\mathbf{d})$. The prior $P(\theta|\mathbf{x}) = P(g(\nu)|\mathbf{x}) = P(\nu|\mathbf{x})$, and the posterior $Q(\theta|\mathbf{d}) = Q(g(\nu)|\mathbf{d}) = Q(\nu|\mathbf{d})$. The first term is the dialogue generation objective in the latent topic inferrer, the

second term is the response generation objective, and the third term is the KL divergence between two Gaussian distributions. All the parameters are learned by optimizing Eq.(14) and updated with back-propagation.

The following previously proposed strategies (Bowman et al., 2016; Zhao et al., 2017) are adopted in training to alleviate the vanishing latent variable problem: (1) KL annealing: the weight of the KL divergence term is gradually increasing from 0 to 1 during training; (2) Bag-of-words loss: the bag-of-words loss requires the latent variable ν , together with the dialogue context, to reconstruct the response bag-of-words representation y_b .

4 Experiments

4.1 Dataset and Competitor Baselines

To ascertain the effectiveness of the proposed model, we construct an open-domain conversation corpus covering a broad range of resources including a movie discussions dataset collected from Reddit (Dodge et al., 2015), an Ubuntu technical corpus (Lowe et al., 2015), and a chit-chat dataset (Zhang et al., 2018). 87,468 context-response pairs were sampled for training, 4,460 for validation and 4,468 for testing.

The code and corpus are available at <http://github.com/hengyicai/AdaND>.

The following state-of-the-art models are adopted as our comparison systems.

SEQ2SEQ The attention-based sequence-to-sequence model (Bahdanau et al., 2015), which is a representative baseline.

CVAE A latent variable conversation model in which it incorporates a latent variable at the sentence-level to inject stochasticity and diversity (Clark and Cao, 2017; Zhao et al., 2017).

LAED A recurrent encoder-decoder conversation model using discrete latent actions for interpretable neural dialogue generation (Zhao et al., 2018).

TA-SEQ2SEQ TA-SEQ2SEQ incorporates the outsourcing topic information into the response generation, where the topics are learned from a separate LDA model to enrich the context (Xing et al., 2017).

DOM-SEQ2SEQ A domain-aware conversation model consisting of multiple domain-targeted

Models	Relevance (%)				Informativeness (%)		
	BLEU	Average	Greedy	Extrema	Distinct-1	Distinct-2	Distinct-3
SEQ2SEQ	0.845	69.60	64.94	45.29	0.2822	0.5922	0.7873
CVAE	1.546	71.23	66.67	47.14	0.5465	1.716	2.731
LAED	0.7545	69.91	63.55	43.12	0.3890	0.9165	1.243
TA-SEQ2SEQ	1.465	72.47	65.9	45.19	0.3593	0.7994	1.016
DOM-SEQ2SEQ	1.189	74.42	66.6	48.47	0.4977	1.294	1.814
ADAND (<i>with context para.</i>)	1.94	74.03	66.76	49.23	0.6493	1.889	2.745
ADAND (<i>with topic para.</i>)	2.051	74.17	66.65	49.04	0.5919	1.699	2.438
ADAND (<i>with both</i>)	1.90	75.59	67.25	51.17	0.7092	2.10	3.108

Table 1: Quantitative evaluation results (%).

SEQ2SEQ models for response generation (Choudhary et al., 2017).

4.2 Evaluation

Following the evaluation procedure in previous work (Li et al., 2016; Xing et al., 2017; Chen et al., 2018), experimental results of all models are reported in terms of the relevance and informativeness. To evaluate the semantic relevance between the generated response and the ground-truth response, we adopted the BLEU metric (Papineni et al., 2002) and three embedding-based similarity metrics proposed in Liu et al. (2016): Embedding Average (**Average**), Embedding Extrema (**Extrema**) and Embedding Greedy (**Greedy**). To measure informativeness and diversity of the response, we exploited the Distinct-1, Distinct-2 and Distinct-3 metrics. A higher ratio of distinct n-grams implies more informative and diverse responses.

4.3 Implementation and Reproducibility

We implemented our model with ParlAI (Miller et al., 2017). The sequence lengths are truncated at 50. We used Adam (Kingma and Ba, 2014) with an initial learning rate of 0.001 to optimize the model. For all the experiments, we employed a 2-layer bidirectional LSTM as the encoder and a unidirectional one as the decoder. The hidden size and the word embedding dimension are both set to 300. The latent variable size is set to 64. The topic number K in our model is set to 5 and the most frequent 3,159 words are taken as the topical words vocabulary by stemming, filtering stop-words from the training set. The batch size is set to 128 for all models. We trained a Twitter LDA model to obtain the topical words for TA-SEQ2SEQ and set its model-specific parameters following the original

paper (Xing et al., 2017). For regularization and preventing over-fitting, a dropout of 0.1 is applied and the weight decay is set to 3×10^{-5} . We used the pretrained word embeddings (Pennington et al., 2014) of 300 dimensions, and the vocabulary size is set to 20,000. All models are trained with early stopping, i.e., if the loss does not decrease after 10 validations. The loss is computed on the validation set at every 0.5 epochs and we save the parameters for the top model on the validation set. We finally report evaluation scores on the test set from the saved model.

4.4 Overall Performance

Table 1 lists the performance of our system and the comparison systems. CVAE and LAED inject SEQ2SEQ with stochastic latent variable, resulting in more informative responses and better performance on Distinct- $\{1, 2, 3\}$. TA-SEQ2SEQ incorporates SEQ2SEQ with the outsourcing topic information from LDA. It is not surprising that it performs much better on the response relevance (BLEU, Average, Greedy, Extrema), while its improvements on the informativeness are limited. DOM-SEQ2SEQ builds multiple domain-specific encoder-decoders. It gains improvements on both the relevance metrics and informativeness metrics.

In general, with both the context-aware and topic-aware parameterization, our model outperforms all the competitive baselines in terms of the response relevance and informativeness.

4.5 Context-aware vs Topic-aware Parameterization

Context-aware parameterization captures local semantics of the given context, while topic-aware parameterization enables parameters sharing among conversations with similar topics. As shown in Ta-

Opponent	Win	Loss	Tie	Kappa
ADAND vs. SEQ2SEQ	62.5%	2.52%	34.98%	46.46
ADAND vs. CVAE	36.98%	7.56%	55.46%	51.59
ADAND vs. LAED	44.54%	6.72%	48.74%	55.57
ADAND vs. TA-SEQ2SEQ	54.62%	5.88%	39.5%	54.91
ADAND vs. DOM-SEQ2SEQ	32.78%	10.08%	57.14%	55.53

Table 2: The results of human evaluation.

Models	Speed (cases/ms)
SEQ2SEQ	0.740
CVAE	0.541
LAED	0.435
TA-SEQ2SEQ	0.172
DOM-SEQ2SEQ	0.106
ADAND	0.407

Table 3: Speed test.

ble 1, both parameterization mechanisms perform much better than the original SEQ2SEQ model, while context-aware parameterization is slightly better in terms of informativeness. When jointly utilizing both the context-aware and topic-aware parameterization mechanisms, we observe the best performance, indicating that these two mechanisms are both beneficial and complementary.

4.6 Human Evaluation

We conducted human evaluations on the test set to further validate the effectiveness of the model. We randomly selected 500 samples from the test set. Three well-educated students were invited to conduct the evaluation. For each case, we provided annotators with triplets (sample, response₁, response₂) whereby one response is generated by ADAND, and the other is generated by a competitor model. The annotators, who have no knowledge about which system the response is from, are then required to independently rate among win (response₁ is better), loss (response₂ is better) and tie (they are equally good or bad), considering four factors: context relevance, logical consistency, fluency and informativeness. Note that if annotators rate different options, this triplet will be counted as “tie”. Table 2 reveals the results of subjective evaluation. The kappa scores indicate that the annotators came to a fair agreement in the judgment.

As expected, ADAND outperforms the other baselines and enjoys a large margin over the existing models. The relative performance of the competitors is consistent with the quantitative evaluation results, confirming the superior performance of our proposed method.

Movie	Politics	Ubuntu	Food	Network
imax	trump	install	food	router
movie	people	grub	vegetarian	wireless
youtube	hillary	apt	seafood	ip
scene	vote	kernel	restaurants	address
marvel	clinton	nvidia	cotto	phone
hulk	election	cd	gourmet	network
avengers	debate	sudo	serves	card
nolan	donald	ssh	starbucks	eth0
comics	support	boot	breakfast	dhcp
batman	working	ubuntu	pizzeria	wifi

Table 4: Topics by the words (β in Eq.(10)) with top-10 highest probability discovered by the latent topic inferer.

4.7 Speed Test

We conducted speed test to verify the efficiency of the ADAND model empirically in Table 3. Augmented with auxiliary components, all the extension models exhibit higher time cost than the original SEQ2SEQ model. We observe that the decoding speeds of CVAE and LAED are relatively comparable with our model. However, when comparing with TA-SEQ2SEQ and DOM-SEQ2SEQ that also elaborately and explicitly model conversations with diverse topics or themes, ADAND shows a clear superiority in decoding speed. For TA-SEQ2SEQ, it relies on an outside LDA model to obtain the topic information. The joint attention and copying mechanism also reduce its efficiency. For DOM-SEQ2SEQ, it is not surprising that the time complexity of multiple topic/theme-specific encoder-decoders is much higher than all-other comparison models. ADAND utilizes a single encoder-decoder and is parameterized dynamically regarding the input context, which ensures its flexibility and efficiency.

4.8 Analysis & Case Study

To get some insights of how topic-aware parameterization performs, we present the topics by the words (β in Eq.(10)) with top-10 highest probabilities in Table 4. The discernible clusters of the topical words (F in Eq.(11)) are illustrated in Figure 4. These evidences demonstrate that the topic inferer in topic-aware parameterization effectively distills the latent topic distribution of each conversation, which enables the parameter sharing among conversations with similar topics.

We also investigate the orthogonality of the learned U and V matrices in Eq.(7) and Eq.(12). We trained our model multiple times with differ-




Figure 4: t-SNE (van der Maaten and Hinton, 2008) projection of topical word embeddings (F in Eq.(11)) in the latent topic inferrer. Words with similar topics are in the same color.

ent parameter initialization methods (drawn values from normal distribution or uniform distribution). We observe that UU^T and VV^T approximate identity matrices. We conjecture that such SVD-alike parameterization implicitly enforces orthogonality during training.

We list several examples generated by different models in Table 5. The inferred latent topic distributions are also presented in the table. It can be observed that responses generated by the original SEQ2SEQ model are more generic. Latent variable conversation models (CVAE and LAED) generate more diverse but sometimes irrelevant responses, TA-SEQ2SEQ tends to produce short responses while DOM-SEQ2SEQ does not perform obviously better than TA-SEQ2SEQ. The responses generated by ADAND are not only relevant but also informative.

5 Related Work

Our work is closely related to the research of dialogue generation in diverse conversations. Previous work relies on external pre-organized topic information (Xing et al., 2017; Wang et al., 2017) or predicted keywords (Yao et al., 2017; Wang et al., 2018) to boost the response informativeness and coherence. Choudhary et al. (2017) further leveraged the topic/theme annotations to build multiple separate encoder-decoder models for topic/theme-aware response generation. In contrast, we do not exploit any outsourcing or labeled topic information. The proposed model directly infers the latent topics of each conversation and is trained in an end-to-end manner. Another difference is that we

	<p>Context: Hi, can you help me find a good movie to watch? Reference: Hi, I would like to help out. SEQ2SEQ: I'm sorry, I'm sorry, I'm sorry. CVAE: I'm looking for a good comedy movie. LAED: I'm sorry, I'm not sure. TA-SEQ2SEQ: I am looking for a good movie. DOM-SEQ2SEQ: I sent the info on your screen, drive carefully! ADAND: What kind of movies do you like?</p> <p>Inferred topic distribution: </p>
2	<p>Context: I love to sit on the couch and play video games myself. Reference: My brother does too, especially since our parents moved to florida recently. SEQ2SEQ: I'm sorry, I'm not a good. CVAE: I'm not sure that is my favorite. LAED: I'm not a good. I'm a good. TA-SEQ2SEQ: That's cool. DOM-SEQ2SEQ: I'm a good. I'm a _unk_ ADAND: That's cool. What do you do for a living?</p> <p>Inferred topic distribution: </p>
3	<p>Context: We go back to celebration in Florida often. It is our home town. Reference: Oh how nice! I love taking my dog to the beach! SEQ2SEQ: I'm sorry, I'm not a good. CVAE: I'm a little more than a lot of time. LAED: I'm sorry, I'm sorry, but there are no results for any other restaurants in the centre. Would you like to try another area? TA-SEQ2SEQ: :) DOM-SEQ2SEQ: I have a 4 star hotel near a park, museum, and theatre, 8.86/10 guest rating, free parking, breakfast, and wifi, from September 2nd to 13th, business class flights, _unk_ 86. ADAND: That's cool. Do you have any hobbies?</p> <p>Inferred topic distribution: </p>
4	<p>Context: Oh, that's all right. No need to reserve my ticket. But I do need the travel time, please. Reference: Your travel time is 60 minutes. Is there anything else that I can do for you? SEQ2SEQ: I have booked you for the reference number is _unk_ CVAE: The travel time is 105 minutes. LAED: I'm sorry, I'm sorry, I'm sorry, but there are no results for that time. Would you like to try another day or time? TA-SEQ2SEQ: The train id is _unk_ DOM-SEQ2SEQ: You're welcome. Have a great day! ADAND: The travel time is 163 minutes. Is there anything else i can help you with?</p> <p>Inferred topic distribution: </p>
5	<p>Context: How is gonna be the weather forecast in San Mateo for the next weekend? Reference: In San Mateo it will be stormy on Saturday, and hail on Sunday. SEQ2SEQ: The address is a park, and the address is a park. CVAE: It will be a low of 60F and a high of 70F. LAED: what city would you like to know the weather for? TA-SEQ2SEQ: Do you have a budget? DOM-SEQ2SEQ: It will not be cloudy this week. ADAND: It will not be cloudy in San Mateo next week.</p> <p>Inferred topic distribution: </p>

Table 5: Test samples of our model (ADAND) and the baselines. The latent topic distributions inferred by ADAND are also presented. The reference is the ground-truth response in the dataset.

maintain a single encoder-decoder for various conversations whereas the model is dynamically and specially parameterized.

The second line of related work is parameterization in NLP. Ha et al. (2017) proposed to train a small network to generate the parameters for another larger network. Such adaptive parameteriza-

tion has been shown to be successful in many NLP tasks, including language modeling (Suarez, 2017; Flennerhag et al., 2018), sequence generation (Ha and Eck, 2018; Peng et al., 2019), and neural machine translation (Platanios et al., 2018). In our work, we parameterize the encoder-decoder with respect to both the context and the latent topics.

Regarding latent variable conversation models, prior researches strive to learn meaningful latent variables for dialogue systems, and reveal that latent variables benefit the neural dialogue models with more diverse response generations (Serban et al., 2017; Zhao et al., 2017; Clark and Cao, 2017; Chen et al., 2018) and interpretable dialogue actions (Zhao et al., 2018). In our model, instead of directly injecting the latent variable into dialogue models, we distill the latent topics through neural variational inference, offering a more interpretable latent variable. Moreover, we parameterize the encoder-decoder with the inferred latent topics, which allows parameter sharing among conversations with similar topics.

6 Conclusion

This paper presents an adaptive neural dialogue generation model—ADAND, which allows the dynamical parameterization of the model to each conversation and enables the generation of appropriate responses in diverse conversations. Specially, we propose two adaptive parameterization approaches: context-aware parameterization which captures local semantics of the input context; and topic-aware parameterization which enables parameter sharing by first inferring the latent topics of the given context and then generating the parameters with the inferred latent topics. The proposed approaches are assessed on a large-scale conversational dataset and the results show that our model achieves superior performance and higher efficiency. It should be noted that our approach is not isolated to only LSTMs. We would like to explore the effectiveness of the approach regarding other structures in future work.

Acknowledgments

We would like to thank all the reviewers for their insightful and valuable comments and suggestions. Hongshen Chen and Cheng Zhang are the corresponding authors.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Luca Bertinetto, João F. Henriques, Jack Valmadre, Philip H. S. Torr, and Andrea Vedaldi. 2016. Learning feed-forward one-shot learners. In *NIPS*, pages 523–531.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Józefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *SIGNLL*, pages 10–21.
- Hongshen Chen, Zhaochun Ren, Jiliang Tang, Yihong Eric Zhao, and Dawei Yin. 2018. Hierarchical variational memory network for dialogue generation. In *WWW*, pages 1653–1662.
- 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 *EMNLP*, pages 1724–1734.
- Sajal Choudhary, Prerna Srivastava, Lyle H. Ungar, and João Sedoc. 2017. [Domain aware neural dialog system](#). *CoRR*, abs/1708.00897.
- Stephen Clark and Kris Cao. 2017. [Latent variable dialogue models and their diversity](#). In *EACL*, pages 182–187.
- Jesse Dodge, Andreea Gane, Xiang Zhang, Antoine Bordes, Sumit Chopra, Alexander H. Miller, Arthur Szlam, and Jason Weston. 2015. [Evaluating prerequisite qualities for learning end-to-end dialog systems](#). *CoRR*, abs/1511.06931.
- Sebastian Flennerhag, Hujun Yin, John A. Keane, and Mark Elliot. 2018. Breaking the activation function bottleneck through adaptive parameterization. In *NeurIPS*, pages 7750–7761.
- David Ha, Andrew M. Dai, and Quoc V. Le. 2017. Hypernetworks. In *ICLR*.
- David Ha and Douglas Eck. 2018. A neural representation of sketch drawings. In *ICLR*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9(8):1735–1780.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *ICLR*.
- Diederik P. Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *NAACL-HLT*, pages 110–119.

- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael 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 *EMNLP*, pages 2122–2132.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. [The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems](#). In *SIGDIAL*, pages 285–294.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *ICML*, volume 70, pages 2410–2419.
- A. H. Miller, W. Feng, A. Fisch, J. Lu, D. Batra, A. Borde, D. Parikh, and J. Weston. 2017. [Parlai: A dialog research software platform](#). *arXiv preprint arXiv:1705.06476*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *ACL*, pages 311–318.
- Hao Peng, Ankur P. Parikh, Manaal Faruqui, Bhuwan Dhingra, and Dipanjan Das. 2019. [Text generation with exemplar-based adaptive decoding](#). In *NAACL-HLT*, pages 2555–2565.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *EMNLP*, pages 1532–1543.
- Emmanouil Antonios Platanios, Mrinmaya Sachan, Graham Neubig, and Tom M. Mitchell. 2018. [Contextual parameter generation for universal neural machine translation](#). In *EMNLP*, pages 425–435.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *AAAI*, 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 *AAAI*, pages 3295–3301.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural responding machine for short-text conversation](#). In *ACL*, pages 1577–1586.
- Xiaoyu Shen, Hui Su, Yanran Li, Wenjie Li, Shuzi Niu, Yang Zhao, Akiko Aizawa, and Guoping Long. 2017. [A conditional variational framework for dialog generation](#). In *ACL*, pages 504–509.
- Joseph Suarez. 2017. Character-level language modeling with recurrent highway hypernetworks. In *NIPS*, pages 3269–3278.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Di Wang, Nebojsa Jojic, Chris Brockett, and Eric Nyberg. 2017. [Steering output style and topic in neural response generation](#). In *EMNLP*, pages 2140–2150.
- Wenjie Wang, Minlie Huang, Xin-Shun Xu, Fumin Shen, and Liqiang Nie. 2018. Chat more: Deepening and widening the chatting topic via A deep model. In *SIGIR*, pages 255–264.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. Topic aware neural response generation. In *AAAI*, pages 3351–3357.
- Lili Yao, Yaoyuan Zhang, Yansong Feng, Dongyan Zhao, and Rui Yan. 2017. [Towards implicit content-introducing for generative short-text conversation systems](#). In *EMNLP*, pages 2190–2199.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *ACL*, pages 2204–2213.
- Tiancheng Zhao, Kyusong Lee, and Maxine Eskenazi. 2018. Unsupervised discrete sentence representation learning for interpretable neural dialog generation. In *ACL*, pages 1098–1107.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. [Learning discourse-level diversity for neural dialog models using conditional variational autoencoders](#). In *ACL*, pages 654–664.