

NEXUS Network: Connecting the Preceding and the Following in Dialogue Generation

Xiaoyu Shen^{1,3*}, Hui Su^{2*}, Wenjie Li⁴ and Dietrich Klakow³

¹Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

²Pattern Recognition Center, Wechat, Tencent, China

³Spoken Language Systems (LSV), Saarland University, Germany

⁴The Hong Kong Polytechnic University, Hong Kong

xshen@mpi-inf.mpg.de, aaronosu@tencent.com

Abstract

Sequence-to-Sequence (seq2seq) models have become overwhelmingly popular in building end-to-end trainable dialogue systems. Though highly efficient in learning the backbone of human-computer communications, they suffer from the problem of strongly favoring short generic responses. In this paper, we argue that a good response should smoothly connect both the preceding dialogue history and the following conversations. We strengthen this connection through mutual information maximization. To sidestep the non-differentiability of discrete natural language tokens, we introduce an auxiliary continuous code space and map such code space to a learnable prior distribution for generation purpose. Experiments on two dialogue datasets validate the effectiveness of our model, where the generated responses are closely related to the dialogue context and lead to more interactive conversations.

1 Introduction

With the availability of massive online conversational data, there has been a surge of interest in building open-domain chatbots with data-driven approaches. Recently, the neural network based sequence-to-sequence (seq2seq) framework (Sutskever et al., 2014; Cho et al., 2014) has been widely adopted. In such a model, the encoder, which is typically a recurrent neural network (RNN), maps the source tokens into a fixed-sized continuous vector, based on which the decoder estimates the probabilities on the target side word by word. The whole model can be efficiently trained by maximum likelihood (MLE) and has demonstrated state-of-the-art performance in various domains. However, this architecture is not

^{*}Indicates equal contribution. X. Shen focuses on algorithm and H. Su is responsible for experiments.

*A*₁: Do you know the movie Star Wars?
*B*₁: Only a bit. **You can tell me about it!**
*A*₂: Of course! This is about ...

Figure 1: A conversation in real life

suitable for modeling dialogues. Recent research has found that while the seq2seq model generates syntactically well-formed responses, they are prone to being off-context, short, and generic. (e.g., “I dont know” or “I am not sure”) (Li et al., 2016a; Serban et al., 2016). The reason lies in the one-to-many alignments in human conversations, where one dialogue context is open to multiple potential responses. When optimizing with the MLE objective, the model tends to have a strong bias towards safe responses as they can be literally paired with arbitrary dialogue context without semantical or grammatical contradictions. These safe responses break the dialogue flow without bringing any useful information and people will easily lose interest in continuing the conversation.

In this paper, we propose NEXUS Network which aims at producing more on-topic responses to maintain an interactive conversation flow. Our assumption is that a good response should serve as a “nexus”: connecting and being informative to both the preceding dialogue context and the follow-up conversations. For example, in Figure 1, the response from *B*₁ is a smooth connection, where the first half indicates the preceding context is a “Do you know” question and the second half informs that the follow-up would be an introduction about *Star Wars*. We establish this connection by maximizing the mutual information (MMI) of the current utterance with both the past and future contexts. In this way, generic responses can be largely discouraged as they contain no valuable information and thus have only weak correlations

with the surrounding context. To enable efficient training, two challenges exist.

The first challenge comes from the discrete nature of language tokens, hindering efficient gradient descent. One strategy is to estimate the gradient by methods like Gumbel-Softmax (Maddison et al., 2017; Jang et al., 2017) or REINFORCE algorithm (Williams, 1992), which has been applied in many NLP tasks (He et al., 2016; Shetty et al., 2017; Gu et al., 2018; Paulus et al., 2018), but the trade-off between bias and variance of the estimated gradient is hard to reconcile. The resulting model usually strongly relies on sensitive hyper-parameter tuning, careful pre-train and task-specific tricks. Li et al. (2016a); Wang et al. (2017) avoid this non-differentiability problem by learning a separate backward model to rerank candidate responses in the testing phase while still adhering to the MLE objective for training. However, the candidate set normally suffers from low diversity and a huge sample size is needed for good performance (Li et al., 2016b).

The second challenge relates to the unknown future context in the testing phase. In our framework, both the history and future context need to be explicitly observed in order to compute the mutual information. When applying it to generating tasks where only the history context is given, there is no way to explicitly take into account the future information. Therefore, reranking-based models do not apply here. (Li et al., 2016c) addresses future information by policy learning, but the model suffers from high variance due to the enormous sequential search space. Serban et al. (2017); Zhao et al. (2017); Shen et al. (2017) adopt the variational inference strategy to reduce the training variance by optimizing over latent continuous variables. However, they all stick to the original MLE objective and no connection with the surrounding context is considered.

In this work, we address both challenges by introducing an auxiliary continuous code space which is learned from the whole dialogue flow. At each time step, instead of directly optimizing discrete utterances, the current, past and future utterances are all trained to maximize the mutual information with this code space. Furthermore, a learnable prior distribution is simultaneously optimized to predict the corresponding code space, enabling efficient sampling in the testing phase without getting access to the ground-truth future con-

versation. Extensive experiments have been conducted to validate the superiority of our framework. The generated responses clearly demonstrate better performance with respect to both coherence and diversity.

2 Model Structure

2.1 Motivation

Let u_i be the i th utterance within a dialogue flow. The dialogue history H_{i-1} contains all the preceding context u_1, u_2, \dots, u_{i-1} and F_{i+1} denotes the future conversations u_{i+1}, \dots, u_T . The objective of our model is to find the decoding probability $p_\theta(u_i|H_{i-1}, F_{i+1})$ that maximizes the mutual information $I(H_{i-1}, u_i)$ and $I(u_i, F_{i+1})$. Formally, the objective is:

$$\begin{aligned} \max_{\theta} \lambda_1 I(H_{i-1}, u_i) + \lambda_2 I(u_i, F_{i+1}) \\ u_i \sim p_\theta(u_i|H_{i-1}, F_{i+1}) \end{aligned} \quad (1)$$

λ_1 and λ_2 adjusts the relative weight. Mutual information is defined over $p_\theta(u_i|H_{i-1}, F_{i+1})$ and the empirical distribution $p(H_{i-1}, F_{i+1})$. Now we assume the future context F_{i+1} is known to us when training the decoding probability, we will address the unknown future problem later.

Directly optimizing with this objective is unfortunately infeasible because the exact computation of mutual information is intractable, and back-propagating through sampled discrete sequences is notoriously difficult to train. The discontinuity prevents the direct application of the reparameterization trick (Kingma and Welling, 2014). Low-variance relaxations like Gumbel-Softmax (Jang et al., 2017), semantic hashing (Kaiser et al., 2018) or vector quantization (van den Oord et al., 2017) lead to biased gradient estimations, which are accumulated as the sequence becomes longer. The Monte-Carlo-Simulation is unbiased but suffers from high variances. Designing a reasonable control variate for variance reduction is an extremely tricky task (Mnih and Gregor, 2014; Tucker et al., 2017). For this sake, we propose replacing u_i with a continuous code space c learned from the whole dialogue flow.

2.2 Continuous Code Space

We define the continuous code space c to follow the Gaussian probability distribution with a diagonal covariance matrix conditioning on the whole

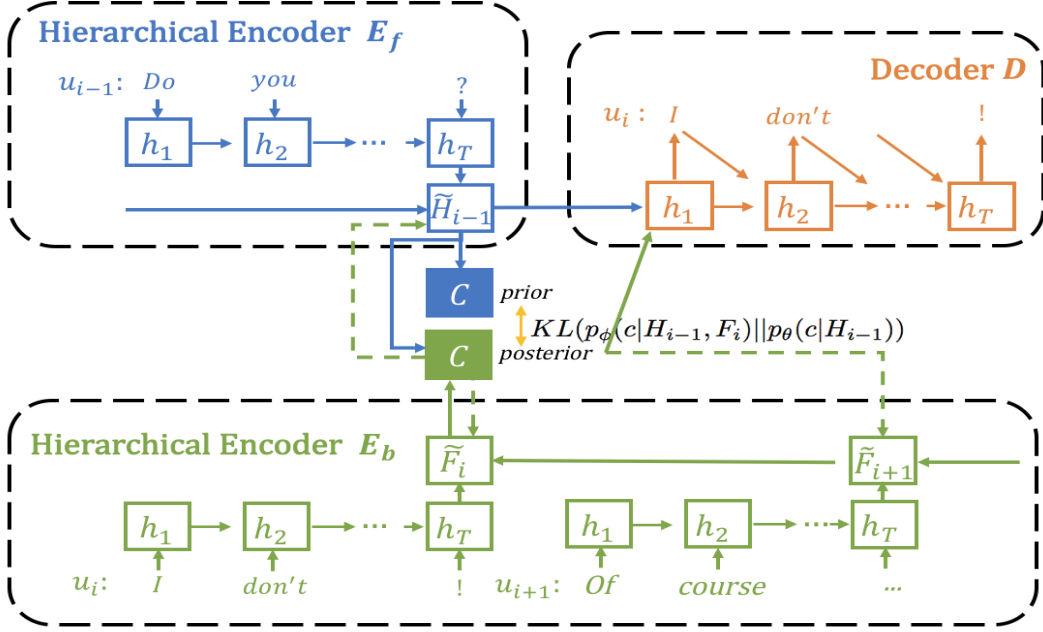


Figure 2: Framework of NEXUS Networks. Full line indicates the generative model to generate the continuous code and corresponding responses. Dashed line indicates the inference model where the posterior code is trained to infer the history, current and future utterances. Both parts are simultaneously trained by gradient descent.

dialogue:

$$c \sim p_\phi(c|H_{i-1}, F_i) = \mathcal{N}(\mu_c, \sigma_c^2 \mathbb{I}|H_{i-1}, F_i) \quad (2)$$

The dialogue history H_{i-1} is encoded into vector \tilde{H}_{i-1} by a forward hierarchical GRU model E_f as in (Serban et al., 2016). The future conversation, including the current utterance, is encoded into \tilde{F}_i by a backward hierarchical GRU E_b . \tilde{H}_{i-1} and \tilde{F}_i are concatenated and a multi-layer perceptron is built on top of them to estimate the Gaussian mean and covariance parameters. The code space is trained to infer the encoded history \tilde{H}_{i-1} and future \tilde{F}_{i+1} . The full optimizing objective is:

$$\begin{aligned} \mathcal{L}(c) &= \max_{\phi} \mathbb{E}_{p_\phi(H_{i-1}, F_i, c)} [\lambda_1 \log p_\phi(\tilde{H}_{i-1}|c) \\ &\quad + \lambda_2 \log p_\phi(\tilde{F}_{i+1}|c)] \\ p_\phi(H_{i-1}, F_i, c) &= p(H_{i-1}, F_i) p_\phi(c|H_{i-1}, F_i) \\ p_\phi(\tilde{H}_{i-1}|c) &= \mathcal{N}(\mu_{H_i}, \sigma_{H_i}^2 \mathbb{I}|c) \\ p_\phi(\tilde{F}_{i+1}|c) &= \mathcal{N}(\mu_{F_{i+1}}, \sigma_{F_{i+1}}^2 \mathbb{I}|c) \end{aligned} \quad (3)$$

where \tilde{H}_{i-1} and \tilde{F}_{i+1} are also assumed to be Gaussian distributed given c with mean and covariance estimated from multi-layer perceptrons. We infer the encoded vectors instead of the original sequences for three reasons. Firstly, inferring dense vectors is parallelizable and computationally much cheaper than autoregressive decoding, especially when the context sequences could

be unlimitedly long. Secondly, sequence vectors can capture more holistic semantic-level similarity than individual tokens. Lastly, It can also help alleviate the posterior collapsing issue (Bowman et al., 2016) when training variational inference models on text (Chen et al., 2017; Shen et al., 2018), which we will use later. It can be shown that the above objective maximizes a lower bound of $\lambda_1 I(H_{i-1}, c) + \lambda_2 I(c, F_{i+1})$, given the conditional probability $p_\phi(c|H_{i-1}, F_i)$. The proof is a direct extension of the derivation in (Chen et al., 2016), followed by the Data Processing Inequality (Beaudry and Renner, 2012) that the encoding function can only reduce the mutual information. As the sampling process contains only Gaussian continuous variables, the above objective can be trained through the reparameterization trick (Kingma and Welling, 2014), which is a low-variance, unbiased gradient estimator (Burda et al., 2015). After training, samples from $p_\phi(c|H_{i-1}, F_i)$ hold high mutual information with both the history and future context. The next step is then transferring the continuous code space to reasonable discrete natural language utterances.

2.3 Decoding from Continuous Space

Our decoder transfers the code space c into the ground-truth utterance u_i by defining the probability distribution $p(u_i|H_{i-1}, c)$, which is imple-

mented as a GRU decoder going through u_i word by word to estimate the output probability. The encoded history H_{i-1} and code space c are concatenated as an extra input at each time step. The loss function for the decoder is then:

$$\begin{aligned}\mathcal{L}(d) &= \max_{\phi} \mathbb{E}_{p_{\phi}(H_{i-1}, F_i, c)} \log p_{\phi}(u_i | H_{i-1}, c) \\ p_{\phi}(H_{i-1}, F_i, c) &= p(H_{i-1}, F_i) p_{\phi}(c | H_{i-1}, F_i)\end{aligned}\quad (4)$$

which can be proved to be the lower bound of the conditional mutual information $I(u_i, c | H_{i-1})$. By maximizing the conditional mutual information, c_i is trained to maintain as much information about the target sequence u_i as possible.

Combining Eq. 3 and 4, our model until now can be viewed as optimizing a lower bound of the following objective:

$$\begin{aligned}\max_{\phi} \lambda_1 I(H_{i-1}, c) + \lambda_2 I(c, F_{i+1}) + I(u_i, c | H_{i-1}) \\ c \sim p_{\phi}(c | H_{i-1}, F_i)\end{aligned}\quad (5)$$

Compared with the original motivation in Eq. 1, we sidestep the non-differentiability problem by replacing u_i with a continuous code space c , then forcing u_i to contain the same information as maintained in c by additionally maximizing the mutual information between them.

Nonetheless, Eq. 5 and Eq. 1 might lead to different optimums as mutual information does not satisfy the transitive law. In the extreme case, different dimensions of c could individually maintain information about history, current and future conversations and the conversations themselves do not share any dependency relation. To avoid this issue, we restrict the dimension of c to be smaller than that of the encoded vectors. In this case, optimizing Eq. 5 will favor utterances having stronger correlations with the surrounding context to achieve a higher total mutual information.

2.4 Learnable Prior Distribution for Unknown Future

The last problem is the sampling mechanism of c in Eq. 2, which conditions on the ground-truth future conversation. In the testing phase, when we have no access to it, we cannot perform the decoding process as in Eq. 4. To allow for decoding with only the history context, we need to learn an appropriate prior distribution $p_{\theta}(c | H_{i-1})$ for c . In

the ideal case, we would like

$$p_{\theta}(c | H_{i-1}) = \sum_{F_i} p_{\phi}(c | H_{i-1}, F_i) = p_{\phi}(c | H_{i-1})\quad (6)$$

However, $p_{\phi}(c | H_{i-1})$ is intractable as it integrates over all possible future conversations. We apply variational inference on c to maximize the variational lower bound (Jordan et al., 1999):

$$\begin{aligned}\mathcal{L}(p) &= \max_{\theta, \phi} \mathbb{E}_{p_{\phi}(c | H_{i-1}, F_i)} \log p_{\theta}(\tilde{F}_i | H_{i-1}, c) \\ &\quad - KL(p_{\phi}(c | H_{i-1}, F_i) || p_{\theta}(c | H_{i-1})) \\ p_{\theta}(\tilde{F}_i | H_{i-1}, c) &\sim \mathcal{N}(\mu_{F_i}, \sigma_{F_i}^2 \mathbb{I} | H_{i-1}, c) \\ p_{\theta}(c | H_{i-1}) &\sim \mathcal{N}(\mu_{prior}, \sigma_{prior}^2 \mathbb{I} | H_{i-1})\end{aligned}\quad (7)$$

It can be reformulated as maximizing:

$$\begin{aligned}\mathbb{E}_{p_{\phi}(c | H_{i-1})} KL(p_{\phi}(\tilde{F}_i | H_{i-1}, c) || p_{\theta}(\tilde{F}_i | H_{i-1}, c)) \\ - KL(p_{\phi}(c | H_{i-1}) || p_{\theta}(c | H_{i-1}))\end{aligned}\quad (8)$$

We can see it implicitly matches $p_{\phi}(c | H_{i-1})$ to a tractable Gaussian distribution $p_{\theta}(c | H_{i-1})$ by minimizing the KL divergence between them. It also functions as a regularizer to prevent overfitting when learning $p_{\phi}(c | H_{i-1}, F_i)$. In the testing phase, we can sample c from the learned prior distribution $p_{\theta}(c | H_{i-1})$, then generate a response based on it.

2.5 Summary

To sum up, the total objective function of our model is:

$$\mathcal{L} = \mathcal{L}(c) + \mathcal{L}(d) + \mathcal{L}(p)\quad (9)$$

Weighting can be added to individual loss functions for better performance, but we find it enough to maintain equal weights and avoid extra hyperparameters. All the parameters are simultaneously updated by gradient descent except for the encoders E_f and E_b , which only accept gradients from $\mathcal{L}(d)$ since otherwise the model can easily learn to encode no information for a lower reconstruction loss in $\mathcal{L}(c)$ and $\mathcal{L}(p)$. An overview of our training procedure is depicted in Fig. 2.

3 Relationship to Existing Methods

MMI decoding MMI decoder was proposed by (Li et al., 2016a) and further extended in (Wang et al., 2017). The basic idea is the same as our model by maximizing the mutual information with

the dialogue context. However, the MMI principle is applied only at the testing phase rather than the training phase. As a result, it can only be used to evaluate the quality of a generation by estimating its mutual information with the context. To apply it in a generative task, we have to first sample some candidate responses with the seq2seq model, then rerank them by accounting for the MMI score. Our model differs from it in that we directly estimate the decoding probability thus no post-sampling rerank is needed. Moreover, we further include the future context to strengthen the connection role of the current utterances.

Conditional Variational Autoencoder The idea of learning an appropriate prior distribution in Eq. 7 is essentially a conditional variational autoencoder (Sohn et al., 2015) where the accumulated posterior distribution is trained to stay close to a prior distribution. It has also been applied in dialogue generation (Serban et al., 2017; Zhao et al., 2017). However, all the above methods stick to the MLE objective function and do not optimize with respect to the mutual information. As we will show in the experiment, they fail to learn the correlation between the utterance and its surrounding context. The generation diversity of these models comes more from the sampling randomness of the prior distribution rather than from the correct understanding of context correlation. Moreover, they suffer from the posterior collapsing problem (Bowman et al., 2016) and require special tricks like KL-annealing, BOW loss or word drop-out (Shen et al., 2018). Our model does not have such problems.

Deep Reinforcement Learning Dialogue Generation (Li et al., 2016c) first considered future success in dialogue generation and applied deep reinforcement learning to encourage more interactive conversations. However, the reward functions are intuitively hand-crafted. The relative weight for each reward needs to be carefully tuned and the training stage is unstable due to the huge search space. In contrast, our model maximizes the mutual information in the continuous space and trains the prior distribution through the reparameterization trick. As a result, our model can be more easily trained with a lower variance. Throughout our experiment, the training process of NEXUS network is rather stable and much less data-hungry. The MMI objective of our model is theoretically

more sound and no manually-defined rules need to be specified.

4 Experiments

4.1 Dataset and Training Details

We run experiments on the DailyDialog (Li et al., 2017b) and Twitter corpus (Ritter et al., 2011). DailyDialog contains 13118 daily conversations under ten different topics. This dataset is crawled from various websites for English learner to practice English in daily life, which is high-quality, less noisy but relatively smaller. In contrast, the Twitter corpus is significantly larger but contains more noise. We obtain the dataset as used in Serban et al. (2017) and filter out tweets that have already been deleted, resulting in about 750,000 multi-turn dialogues. The contents have more informal, colloquial expressions which makes the generation task harder. These two datasets are randomly separated into training/validation/test sets with the ratio of 10:1:1.

In order to keep our model comparable with the state-of-the-art, we keep most parameter values the same as in (Serban et al., 2017). We build our vocabulary dictionary based on the most frequent 20,000 words for both corpus and map other words to a UNK token. The dimensionality of the code space c is 100. We use a learning rate of 0.001 for DailyDialog and 0.0002 for Twitter corpus. The batch size is fixed to 128. The word vector dimension is 300 and is initialized with the public Word2Vec (Mikolov et al., 2013) embeddings trained on the Google News Corpus. The probability estimators for the Gaussian distributions are implemented as 3-layer perceptrons with the hyperbolic tangent activation function. As mentioned above, when training NEXUS models, we block the gradient from $\mathcal{L}(c)$ and $\mathcal{L}(p)$ with respect to E_f and E_b to encourage more meaningful encodings. The UNK token is prevented from being generated in the test phase. We implemented all the models with the open-sourced Python library Pytorch (Paszke et al., 2017) and optimized using the Adam optimizer (Kingma and Ba, 2015).

4.2 Compared Models

We conduct extensive experiments to compare our model against several representative baselines.

Seq2Seq: Following the same implementation as in (Vinyals and Le, 2015), the seq2seq model serves as a baseline. We try both greedy decoding

Model	DailyDialog			Twitter		
	Average	Greedy	Extreme	Average	Greedy	Extreme
Greedy	0.443	0.376	0.328	0.510	0.341	0.356
Beam	0.437	0.350	0.369	0.505	0.345	0.352
MMI	0.457	0.371	0.371	0.518	0.353	0.365
RL	0.405	0.329	0.305	0.460	0.349	0.323
VHRED	0.491	0.375	0.313	0.525	0.389	0.372
NEXUS-H	0.479	0.381*	0.385*	0.558*	0.392	0.373
NEXUS-F	0.476	0.383*	0.373	0.549*	0.393	0.386*
NEXUS	0.488	0.392*	0.384*	0.556*	0.397*	0.391*

Table 1: Results of embedding-based metrics. * indicates statistically significant difference ($p < 0.05$) from the best baselines. The same mark is used in Table 2

and beam search (Graves, 2012) with beam size set to 5 when testing.

MMI: We implemented the bidirectional-MMI decoder as in Li et al. (2016a), which showed better performance over the anti-LM model. The hyperparameter λ is set to 0.5 as suggested. 200 candidates per context are sampled for re-ranking.

VHRED: The VHRED model is essentially a conditional variational autoencoder with hierarchical encoders (Serban et al., 2017; Zhao et al., 2017). To alleviate the posterior collapsing problem, we apply the KL-annealing trick and early stop with the step set as 12,000 for the DailyDialog and 75,000 for the Twitter corpus.

RL: Deep reinforcement learning chatbot as in (Li et al., 2016c). We use all the three reward functions mentioned in the paper and keep the relative weights the same as in the original paper. Policy network is initialized with the above-mentioned MMI model.

NEXUS-H: NEXUS network maximizing mutual information only with the history ($\lambda_2 = 0$).

NEXUS-F: NEXUS network maximizing mutual information only with the future ($\lambda_1 = 0$).

NEXUS: NEXUS network maximizing mutual information with both the history and future.

NEXUS-H and NEXUS-F are implemented to help us better analyze the effects of different components in our model. The hyperparameters λ_1 and λ_2 in NEXUS are set to be 0.5 and 1 respectively as we find history vector is consistently easier to be reconstructed than the future vector (A.6).

4.3 Metric-based Performance

Embedding Score We conducted three embedding-based evaluations (average, greedy

and extrema) (Liu et al., 2016), which map responses into vector space and compute the cosine similarity (Rus and Lintean, 2012). The embedding-based metrics can to a large extent capture the semantic-level similarity between generated responses and ground truth. We represent words using Word2Vec embeddings trained on the Google News Corpus. We also measure the uncertainty of the score by assuming each data point is independently Gaussian distributed. The standard deviation yields the 95% confidence interval (Barany et al., 2007). Table 1 reports the embedding scores on both datasets. NEXUS network significantly outperforms the best baseline model in most cases. Notably, NEXUS can absorb the advantages from both NEXUS-H and NEXUS-F. The history and future information seem to help the model from different perspectives. Taking into account both of them does not create a conflict and the combination leads to an overall improvement. RL performs rather poorly on this metric, which is understandable as it does not target the ground-truth responses during training (Li et al., 2016c).

BLEU Score BLEU is a popular metric that measures the geometric mean of the modified n-gram precision with a length penalty (Papineni et al., 2002). Table 2 reports the BLEU 1-3 scores. Compared with embedding-based metrics, the BLEU score quantifies the word-overlap between generated responses and the ground-truth. One challenge of evaluating dialogue generation by BLEU score is the difficulty of accessing multiple references for the one-to-many alignment relation. Following Sordani et al. (2015); Zhao et al.

Model	DailyDialog			Twitter		
	BLEU-1	BLEU-2	BLEU-3	BLEU-1	BLEU-2	BLEU-3
Greedy	0.394	0.245	0.157	0.340	0.203	0.116
Beam	0.386	0.251	0.163	0.338	0.205	0.112
MMI	0.407	0.269	0.172	0.347	0.208	0.118
RL	0.298	0.186	0.075	0.314	0.199	0.103
VHRED	0.395	0.281	0.190	0.355	0.211	0.124
NEXUS-H	0.418	0.279	0.199*	0.366*	0.212	0.126
NEXUS-F	0.399	0.260	0.167	0.359	0.213	0.123
NEXUS	0.424*	0.276	0.198*	0.363*	0.220*	0.131*

Table 2: Results of BLEU score. It is computed based on the smooth BLEU algorithm (Lin and Och, 2004). p-value interval is computed base on the altered bootstrap resampling algorithm (Riezler and Maxwell, 2005)

(2017); Shen et al. (2018), for each context, 10 more candidate references are acquired by using information retrieval methods (see Appendix A.4 for more details). All candidates are then passed to human annotators to filter unsuitable ones, resulting in 6.74 and 5.13 references for DailyDialog and Twitter dataset respectively. The human annotation is costly, so we evaluate it on 1000 sampled test cases for each dataset. As the BLEU score is not the simple mean of individual sentence scores, we compute the 95% significance interval by bootstrap resampling (Koehn, 2004; Riezler and Maxwell, 2005). As can be seen, NEXUS network achieves best or near-best performances with only greedy decoders. NEXUS-H generally outperforms NEXUS-F as the connection with future context is not explicitly addressed by the BLEU score metric. MMI and VHRED bring minor improvements over the seq2seq model. Even when evaluated on multiple references, RL still performs worse than most models.

Connecting the preceding We define two metrics to evaluate the model’s capability of “connecting the preceding context”: **AdverSuc** and **Neg-PMI**. AdverSuc measures the coherence of generated responses with the provided context by learning an adversarial discriminator (Li et al., 2017a) on the same corpus to distinguish coherent responses from randomly sampled ones. We encode the context and response separately with two different LSTM neural networks and output a binary signal indicating coherent or not¹. The Adver-

¹We apply the same architecture as in Lu et al. (2017). In our experiment, the discriminator performs reasonably well in the 4 scenarios outlined in Li et al. (2017a) and thus can be used as a fair evaluation metric.

Suc value is reported as the success rate that the model fools the classifier into believing its false generations ($p(\text{generated} = \text{coherent}) > 0.5$). Neg-PMI measures the negative pointwise mutual information value $-\log p(c|r)/p(c)$ between the generated response r and the dialogue context c . $p(c|r)$ is estimated by training a separate backward seq2seq model. As $p(c)$ is a constant, we ignore it and only report the value of $-\log p(c|r)$. A good model should achieve a higher Adver-Suc and a lower Neg-PMI. The results are listed in Table 3. We can see there is still a big gap between ground-truth and synthesized responses. As expected, NEXUS-H leads to the most significant improvement. MMI model also performs remarkably well, but it requires post-reranking thus the sampling process is much slower. VHRED and NEXUS-F do not help much here, sometimes even slightly degrade the performance. We also tried removing the history context when computing the posterior distribution in VHRED, the resulting model has similar performance among all metrics, which suggests VHRED itself cannot actually learn the correlation pattern with the preceding context. Surprisingly, though RL explicitly set the coherence score as a reward function, its performance is far from satisfying. We assume RL requires much more data to learn the appropriate policy than other models and the training process suffers from a higher variance. The result is thus hard to be guaranteed.

Connecting the following We measure the model’s capability of “connecting the following context” from two perspectives: number of the simulated turns and diversity of generated re-

Model	AdverSuc	Neg-PMI	#Turns	Distinct-1	Distinct-2	Pri	Post	Flu
Greedy	0.21 0.13	47.4 45.8	0.2 0.6	.019 .017	.096 .072	0.45	0.04	0.92
Beam	0.16 0.12	47.2 45.3	0.2 0.7	.026 .019	.103 .086	0.52	0.06	0.90
MMI	0.30 0.19	45.6 43.2	1.1 1.6	.042 .025	.247 .117	0.56	0.13	0.89
RL	0.13 0.11	45.0 42.6	2.3 2.3	.048 .033	.324 .287	0.46	0.15	0.69
VHRED	0.19 0.16	46.8 44.7	1.7 1.1	.255 .106	.431 .311	0.42	0.22	0.92
NEXUS-H	0.36 0.21	44.1 41.8	2.0 1.8	.263 .108	.454 .306	0.66	0.20	0.92
NEXUS-F	0.22 0.12	47.1 45.9	2.6 2.2	.288 .117	.466 .325	0.51	0.31	0.94
NEXUS	0.35 0.18	44.6 41.4	2.8 2.5	.282 .119	.470 .329	0.70	0.33	0.93
GROUND	0.87 0.73	40.5 38.1	4.8 4.0	.390 .215	.522 .495	0.92	0.67	0.97

Table 3: Coherence, diversity and human evaluations. Left: DailyDialog results, right: Twitter results

sponses. We apply all models to generate multiple turns until a generic response is reached. The set of generic responses is manually examined to include all utterances providing only passive dull replies². The number of generated turns can reflect the time that a model can maintain an interactive conversation. The results are reflected in the **#Turns** column in Table 3. As in (Li et al., 2016a), we measure the diversity by the percentage of distinct unigrams (**Distinct-1**) and bigrams (**Distinct-2**) in all generated responses. Intuitively a higher score on these three metrics implies a more interactive generation system that can better connect the future context. Again, NEXUS network dominates most fields. NEXUS-F brings more impact than NEXUS-H as it explicitly encourages more interactive turns. Most seq2seq models fail to provide an informative response in the first turn. The MMI-decoder does not change much, possibly because the sampling space is not large enough, a more diverse sampling mechanism (Vijayakumar et al., 2018) might help. NEXUS network can effectively continue the conversation for 2.8 turns for DailyDialog and 2.5 turns for Twitter, which is closest to the ground truth (4.8 and 4.0 turns respectively). It also achieves the best diversity score in both datasets. It is worth mentioning that NEXUS-H also improves over baselines, though not as significantly as NEXUS-F, so NEXUS is not a trade-off but more like an enhanced version from NEXUS-H and NEXUS-F.

In summary, NEXUS network clearly generates higher-quality responses in both coherence and diversity, even in a rather small dataset like DailyDialog. NEXUS-H contributes more to the coher-

²We use a simple rule matching method (see Appendix A.5). We manually inspect it on a validation subset and find the accuracy is more than 90%. Similar methods are adopted in (Li et al., 2016c).

ence and NEXUS-F more to the diversity.

4.4 Human Evaluation

We also employed crowdsourced judges to provide evaluations for a random sample of 500 items in the DailyDialog test dataset. Participants are asked to assign a binary score to each context-response pair from three perspectives: whether the response coincides with its preceding context (Pri), whether the response is interesting enough for people to continue (Post) and whether the response itself is a fluent natural sentence (Flu). Each sample gets one point if judged as yes and zero otherwise. Each pair is judged by three participants and the score supported by most people is adopted. We also evaluated the inter-annotator consistency by Fleiss’k score (Fleiss, 1971) and obtained k scores of 0.452 for Pri, 0.459 for Post (moderate agreement) and 0.621 for Flu (substantial agreement), which implies most context-response pairs reach a consensus on the evaluation task. We compute the average human score for each model. Unlike metric-based scores, the human evaluation is conducted only on the DailyDialog corpus as it contains less noise and can be more fairly evaluated by human judges. Table 3 shows the result in the last three columns. As can be seen, the pri and post human scores are highly correlated with the automatic evaluation metric “coherence” and “#turns”, verifying the validity of these two metrics. As for fluency, there is no significant difference among most models. As we also manually examined, fluency is not a major problem and all models produce mostly well-formed sentences. Overall, NEXUS network does produce responses that are more acceptable to human judges.

Table 4 presents some randomly sampled context-response pairs provided by MMI,

Context	Response
Hi , Alice , what are you going to do this sunday ? __eou__	MMI: Oh , that' s great ! VHRED : Well , it seems that i' m going to be late . RL : How long he it be ? NEXUS: Nothing really , what ' s up ?
Did you watch the Oscars on Sunday night ? __eou__	MMI: I have no idea . VHRED: That' s not too bad . RL : It' s my mum' s birthday on Sunday . NEXUS: No , nothing much. Could you help me ?
Let' s go for a drink ! __eou__	MMI: Thanks . VHRED: all right . let me see you then . RL : Thank you inviting me . NEXUS: Sounds good . Where are we going ?
You won' t want to miss our webinar tomorrow ! __eou__	MMI: lol ! VHRED: That ' s terrific ! RL : Sure, where where ? NEXUS: Thanks for your invitation! I' m free !
How I read the article ? __eou__	MMI: there!!! VHRED: What ' s good with it ? RL : : I dont know what to do with it ! NEXUS: Maybe the force is yourself !

Table 4: Examples of context-response pairs. __eou__ denotes end-of-utterance. First three rows are from DailyDialog and the last two rows are from Twitter

VHRED, RL and NEXUS model. We see NEXUS network does generate more interactive outputs than the other three. Though reranked by the bidirectional language model, the MMI decoder still produces quite a few generic responses. VHRED's utterances are more diverse, but it only cares about answering to the immediate query and makes no efforts to bring about further topics. Moreover, it also generates more inappropriate responses than the others. RL provides diverse responses but sometimes not fluent or coherent enough. We do observe that NEXUS sometimes generate over-complex questions which are not very natural, as in the second example. But in most cases, it outperforms the others.

5 Conclusion

In this paper, we propose "NEXUS Network" to enable more interactive human-computer conversations. The main goal of our model is to strengthen the "nexus" role of the current utterance, connecting both the preceding and the following dialogue context. We compare our model with MMI, reinforcement learning and CVAE-based models. Experiments show that NEXUS network consistently produces higher-quality re-

sponses. The model is easier to train, requires no special tricks and demonstrates remarkable generalization capability even in a very small dataset.

Our model can be considered as combining the objective of MMI and CVAE and is compatible with current improving techniques. For example, mutual information can be maximized under a tighter bound using Donsker-Varadhan or f-divergence representation (Donsker and Varadhan, 1983; Nowozin et al., 2016; Belghazi et al., 2018). Extending the code space distribution to more than Gaussian by importance weighted autoencoder (Burda et al., 2015), inverse autoregressive flow (Kingma et al., 2016) or VampPrior (Tomczak and Welling, 2018) should also help with the performance.

Acknowledgments

We thank all anonymous reviewers, Gerhard Weikum, Jie Zhou, Cheng Niu and the dialogue system team of Wechat AI for valuable comments. Xiaoyu Shen is supported by IMPRS-CS fellowship. This work is partially funded by DFG collaborative research center SFB 1102 and Research Grants Council of Hong Kong (PolyU 152036/17E, 152040/18E).

References

- Imre Barany, Van Vu, et al. 2007. Central limit theorems for gaussian polytopes. *The Annals of Probability*, 35(4):1593–1621.
- Normand J Beaudry and Renato Renner. 2012. An intuitive proof of the data processing inequality. *Quantum Information & Computation*, 12(5-6):432–441.
- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Devon Hjelm, and Aaron Courville. 2018. Mutual information neural estimation. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 531–540, Stockholmsmssan, Stockholm Sweden. PMLR.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. 2015. Importance weighted autoencoders. *CoRR*, abs/1509.00519.
- Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2172–2180.
- Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2017. Variational lossy autoencoder. *ICLR*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734.
- Monroe D Donsker and SR Srinivasa Varadhan. 1983. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on Pure and Applied Mathematics*, 36(2):183–212.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Alex Graves. 2012. Sequence transduction with recurrent neural networks. *CoRR*, abs/1211.3711.
- Jiatao Gu, Daniel Jiwoong Im, and Victor OK Li. 2018. Neural machine translation with gumbel-greedy decoding. *AAAI*, pages 5125–5132.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tieyan Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *Advances in Neural Information Processing Systems*, pages 820–828.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. *ICLR*.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. 1999. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233.
- Lukasz Kaiser, Samy Bengio, Aurko Roy, Ashish Vaswani, Niki Parmar, Jakob Uszkoreit, and Noam Shazeer. 2018. Fast decoding in sequence models using discrete latent variables. In *Proceedings of the 35th International Conference on Machine Learning Research*, pages 2390–2399, Stockholmsmssan, Stockholm Sweden. PMLR.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *ICLR*.
- Diederik P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. 2016. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pages 4743–4751.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. *ICLR*.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 conference on empirical methods in natural language processing*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. A simple, fast diverse decoding algorithm for neural generation. *CoRR*, abs/1611.08562.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016c. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017a. Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2157–2169.

- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017b. Dailydialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 986–995.
- Chin-Yew Lin and Franz Josef Och. 2004. Orange: a method for evaluating automatic evaluation metrics for machine translation. In *Proceedings of the 20th international conference on Computational Linguistics*, page 501. Association for Computational Linguistics.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132.
- Yichao Lu, Phillip Keung, Shaonan Zhang, Jason Sun, and Vikas Bhardwaj. 2017. A practical approach to dialogue response generation in closed domains. *CoRR*, abs/1703.09439.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The concrete distribution: A continuous relaxation of discrete random variables. *ICLR*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *ICLR workshop*.
- Andriy Mnih and Karol Gregor. 2014. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, pages 1791–1799.
- Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. 2016. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pages 271–279.
- Aaron van den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pages 6306–6315.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS workshop*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. *ICLR*.
- Stefan Riezler and John T Maxwell. 2005. On some pitfalls in automatic evaluation and significance testing for mt. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 57–64.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. Data-driven response generation in social media. In *Proceedings of the conference on empirical methods in natural language processing*, pages 583–593. Association for Computational Linguistics.
- Vasile Rus and Mihai Lintean. 2012. A comparison of greedy and optimal assessment of natural language student input using word-to-word similarity metrics. In *Proceedings of the Seventh Workshop on Building Educational Applications Using NLP*, pages 157–162. Association for Computational Linguistics.
- Iulian V Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 3776–3783. AAAI Press.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*, pages 3295–3301.
- 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 *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 504–509.
- Xiaoyu Shen, Hui Su, Shuzi Niu, and Vera Demberg. 2018. Improving variational encoder-decoders in dialogue generation. *AAAI*, pages 5456–5463.
- Rakshith Shetty, Marcus Rohrbach, Lisa Anne Hendricks, Mario Fritz, and Bernt Schiele. 2017. Speaking the same language: Matching machine to human captions by adversarial training. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, pages 3483–3491.
- Alessandro Sordoni, 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 *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational*

- Linguistics: Human Language Technologies*, pages 196–205.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Jakub Tomczak and Max Welling. 2018. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pages 1214–1223.
- George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. 2017. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pages 2627–2636.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasaath R Selvaraju, Qing Sun, Stefan Lee, David J Crandall, and Dhruv Batra. 2018. Diverse beam search for improved description of complex scenes. *AAAI*, pages 7371–7379.
- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. *CoRR*, abs/1506.05869.
- Di Wang, Nebojsa Jojic, Chris Brockett, and Eric Nyberg. 2017. Steering output style and topic in neural response generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2140–2150.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 654–664.