

Better Conversations by Modeling, Filtering, and Optimizing for Coherence and Diversity

Xinnuo Xu, Ondřej Dušek, Ioannis Konstas and Verena Rieser
The Interaction Lab, School of Mathematical and Computer Sciences
Heriot-Watt University, Edinburgh
xx6, o.dusek, i.konstas, v.t.rieser@hw.ac.uk

Abstract

We present three enhancements to existing encoder-decoder models for open-domain conversational agents, aimed at effectively modeling coherence and promoting output diversity: (1) We introduce a measure of coherence as the GloVe embedding similarity between the dialogue context and the generated response, (2) we filter our training corpora based on the measure of coherence to obtain topically coherent and lexically diverse context-response pairs, (3) we then train a response generator using a conditional variational autoencoder model that incorporates the measure of coherence as a latent variable and uses a context gate to guarantee topical consistency with the context and promote lexical diversity. Experiments on the OpenSubtitles corpus show a substantial improvement over competitive neural models in terms of BLEU score as well as metrics of coherence and diversity.

1 Introduction

End-to-end neural response generation methods are promising for developing open domain dialogue systems as they allow to learn from very large unlabeled datasets (Shang et al., 2015; Sordoni et al., 2015; Vinyals and Le, 2015). However, these models have also been shown to generate generic, uninformative, and non-coherent replies (e.g., “*I don’t know.*” in Figure 1), mainly due to the fact that neural systems tend to settle for the most frequent options, thus penalizing length and favoring high-frequency word sequences (Soutsov and Sarawagi, 2016; Wei et al., 2017).

To address these problems, Li et al. (2016a) and Li et al. (2017a) attempt to promote diversity by improving the objective function, but do not model diversity explicitly. Serban et al. (2017) focus on

Conversational history	Response
A: You stay out of this. B: So you want <u>water</u> , huh? A: That’s right.	B-Coh: Well, I got <u>water</u> . B-Incoh: I don’t know.
A: Where do we start? B: <u>Kitchen</u> . A: Definitely the <u>kitchen</u> .	B-Coh: Specifically the <u>stove</u> . B-Incoh: Let’s go for a walk.

Figure 1: Examples of conversational history (left) with two alternative responses to follow it (right): (B-Coh) a more coherent, topical utterance, and (B-Incoh) a generic, inconsistent response.

model structure without any upgrades to the objective function. Other works control the style of the output by leveraging external resources (Hu et al. (2017): sentiment classifier, time annotation; Zhao et al. (2017): dialogue acts) or focus on well-structured input such as paragraphs (Li and Jurafsky, 2017).

This paper extends previous attempts to model diversity and coherence by enhancing all three aspects of the learning process: the data, the model, and the objective function. While previous research has addressed these aspects individually, this paper is the first to address all three in a unified framework. Instead of using existing linguistic knowledge or labeled datasets, we aim to control for coherence by learning directly from data, using a fully unsupervised approach. This is also the first work encoding and evaluating coherence explicitly in the dialogue generation task, as opposed to using diversity, style, or other properties of responses as a proxy.

In this work, given a dialogue history, we regard as a coherent response an utterance that is thematically correlated and naturally continuing from the previous turns, as well as lexically diverse. For example, in Figure 1 the response “*Specifically the stove.*” is a very natural and coherent response, elaborating on the topic of *kitchen* introduced in

the previous two utterances and containing rich thematic words, whereas the response “*Let’s go for a walk.*” is unrelated and uninteresting.

In order to obtain coherent responses, we present three generic enhancements to existing encoder-decoder (E-D) models:

1. We define a *measure of coherence* simply as the averaged word embedding similarity between the words of the context and the response computed using GloVe vectors (Pennington et al., 2014).
2. We filter a corpus of conversations based on our measure of coherence, which leaves us with context-response pairs that are both topically coherent and lexically diverse.
3. We train an E-D generator recast as a conditional Variational Autoencoder (cVAE; Zhao et al., 2017) model that incorporates two latent variables, one for encoding the context and another for conditioning on the measure of coherence, trained jointly as in Hu et al. (2017). We then decode using a context gate (Tu et al., 2017) to control the generation of words that directly relate to the most topical words of the context and promote coherence.

Experiments on the OpenSubtitles (Lison and Meena, 2016) corpus demonstrate the effectiveness of the overall approach. Our models achieve a substantial improvement over competitive neural models. We provide an ablation analysis, quantifying the contributions that come from effective modeling of coherence into our models. All our experimental code is freely available on GitHub.¹

2 Coherence-based Dialogue Generation

Our model aims to generate responses given a dialogue context, incorporating measures of coherence estimated purely from the training data. We propose the following enhancements to the attention-based E-D architecture (Bahdanau et al., 2015; Luong et al., 2015):

- We introduce a stochastic latent variable \mathbf{z} conditioned on previous dialogue context to store the global information about the conversation (Bowman et al., 2016; Chung et al., 2015; Li and Jurafsky, 2017; Hu et al., 2017).

¹https://github.com/XinnuoXu/CVAE_Dial

- We force the model to condition on the measure of coherence explicitly by encoding a latent variable (code) c learned from data.
- We incorporate a context gate (Tu et al., 2017) that dynamically controls the ratio at which the generated words in the response derive directly from the coherence-enhanced dialogue context or the previously generated parts of the response.

In the rest of this section, we introduce the measure of coherence (Section 2.1), we present an overview of our model (Section 2.2), and finally describe the model in detail (Sections 2.3–2.4).

2.1 Measure of Dialogue Coherence

Semantic vector space models of language represent each word with a real-valued word embedding vector (Pennington et al., 2014). By simply taking a weighted average of all its word embeddings, a whole sentence can be mapped into the semantic vector space. We define the coherence of a dialogue as the average distance between semantic vectors of preceding dialogue context and its response.

Let $\mathbf{x} = \{x_1, \dots, x_j, \dots, x_J\}$ represent a dialogue context and $\mathbf{y} = \{y_1, \dots, y_i, \dots, y_I\}$ a response. J and I are the numbers of words in the dialogue context and its response, respectively. Semantic vector space models map each word x_j into embeddings \mathbf{x}_j^{emb} , and y_i into \mathbf{y}_i^{emb} . The semantic representation of a dialogue context \mathbf{x} is then $\mathbf{x}^{emb} = \sum_{j=1}^J w_j \mathbf{x}_j^{emb}$; for a response \mathbf{y} , it is $\mathbf{y}^{emb} = \sum_{i=1}^I v_i \mathbf{y}_i^{emb}$. Here, w_j and v_i are importance weights for each word in the sentence.² The measure of coherence is then defined as the cosine distance of the two semantic vectors of the dialogue context and its response:

$$C(\mathbf{x}, \mathbf{y}) = \cos(\mathbf{x}^{emb}, \mathbf{y}^{emb}) \quad (1)$$

2.2 Model Overview

End-to-end response generation for dialogue can be formalized as follows: Given a dialogue context \mathbf{x} , a dialogue generator generates the next utterance \mathbf{y} . During the training process, the aim for a dialogue generator is to maximize the probability $p(\mathbf{y}|\mathbf{x})$ over the training dataset. To encode

²We set the importance weights to 0 for a list of stop words (high-frequency words such as articles and prepositions, names, punctuation marks), 1 otherwise.

dialogue contexts that adequately incorporate coherence information, we build our generator based on the cVAE model of Hu et al. (2017), which has been used to control text generation with respect to linguistic properties, such as tense or sentiment.

In our model, the response \mathbf{y} is generated conditioned on the previous conversation \mathbf{x} , a diversity-promoting latent variable \mathbf{z} , and a latent variable c indicating dialogue coherence; \mathbf{z} and c are independent. The generation probability $p(\mathbf{y}|\mathbf{x})$ is defined as:

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}) &= \int_{\mathbf{z},c} p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c) p(\mathbf{z}, c|\mathbf{x}) d\mathbf{z} dc \\ &= \int_{\mathbf{z},c} p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c) p(\mathbf{z}|\mathbf{x}) p(c|\mathbf{x}) d\mathbf{z} dc \end{aligned} \quad (2)$$

Unfortunately, optimizing Eq (2) during training is intractable; therefore, we apply variational inference and optimize instead the variational lower bound:

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x}) &= \log \int_{\mathbf{z},c} p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c) p(\mathbf{z}, c|\mathbf{x}) d\mathbf{z} dc \\ &\geq \mathbb{E}_{q(\mathbf{z}|\mathbf{x}, \mathbf{y})p(c|\mathbf{x}, \mathbf{y})} [\log p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c)] \\ &\quad - D_{KL}(q(\mathbf{z}|\mathbf{x}, \mathbf{y}) \| p(\mathbf{z}|\mathbf{x})) \end{aligned} \quad (3)$$

where $p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c)$ is the probability of generating utterance \mathbf{y} given \mathbf{x} , \mathbf{z} and c ; $q(\mathbf{z}|\mathbf{x}, \mathbf{y})$ stands for the approximate posterior distribution of the latent variable \mathbf{z} conditioned on dialogue context \mathbf{x} and the gold response \mathbf{y} ; $p(c|\mathbf{x}, \mathbf{y})$ is the measure of coherence between context \mathbf{x} and response \mathbf{y} ; $p(\mathbf{z}|\mathbf{x})$ is the true prior distribution of \mathbf{z} conditioned only on dialogue context \mathbf{x} ; $D_{KL}(\cdot|\cdot)$ denotes the KL-divergence. We assume that both $q(\mathbf{z}|\mathbf{x}, \mathbf{y})$ and $p(\mathbf{z}|\mathbf{x})$ are Gaussian with mean vectors μ_{appr} , μ_{true} and covariance matrices Σ_{appr} , Σ_{true} .

2.2.1 Model Details

Optimizing Eq (3) consists of two parts: (1) minimizing the KL-divergence between the approximate posterior distribution and the true prior distribution of \mathbf{z} , (2) maximizing the probability of generating the gold response \mathbf{y} conditioned on dialogue context \mathbf{x} and coherence factors \mathbf{z} and c . Figure 2 shows the pipeline of the training procedure.

Encoder: First, we encode a dialogue context \mathbf{x} into a hidden state \mathbf{h} using the *context encoder*, which is based on Recurrent Neural Networks (RNNs). Then the *posterior network* encodes both dialogue context \mathbf{x} and gold response \mathbf{y} into a hidden state \mathbf{h}_{appr} followed by two linear transformations $f_{appr}(\cdot)$ and $g_{appr}(\cdot)$ to map \mathbf{h}_{appr} into mean vector μ_{appr} and covariance matrix Σ_{appr} . The latent variable \mathbf{z} can be sampled from the distribution $N(\mu_{appr}, \Sigma_{appr})$:

$$\begin{aligned} \mu_{appr} &= f_{appr}(\mathbf{h}_{appr}) \\ \Sigma_{appr} &= g_{appr}(\mathbf{h}_{appr}) \\ q(\mathbf{z}|\mathbf{x}, \mathbf{y}) &= N(\mu_{appr}, \Sigma_{appr}) \end{aligned} \quad (4)$$

The *prior network* in Figure 2 takes a form similar to the posterior network:

$$\begin{aligned} \mu_{true} &= f_{true}(\mathbf{h}_{true}) \\ \Sigma_{true} &= g_{true}(\mathbf{h}_{true}) \\ p(\mathbf{z}|\mathbf{x}) &= N(\mu_{true}, \Sigma_{true}) \end{aligned} \quad (5)$$

where \mathbf{h}_{true} is the final hidden state of an RNN encoding only the dialogue context \mathbf{x} , and $f_{true}(\cdot)$, $g_{true}(\cdot)$ are linear transformations. Code c is given by the *coherence measure* from Eq (1).

Decoder: We build an attention-based decoder (Bahdanau et al., 2015; Luong et al., 2015) using RNNs to generate responses conditioned on encoded dialogue context \mathbf{h} , diversity signal \mathbf{z} , and coherence signal c . We concatenate the latent variables \mathbf{z} and c to the context encoder hidden state \mathbf{h} and feed them into the decoder as the initial hidden state \mathbf{s}_0 , similar to Hu et al. (2017).

During the decoding process, tokens are generated sequentially under the following probability distribution:

$$\begin{aligned} p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c) &= \prod_{i=1}^I p(y_i|y^{<i}, \mathbf{x}, \mathbf{z}, c) \\ &= \prod_{i=1}^I g(y_{i-1}, \mathbf{s}_i, \mathbf{a}_i) \end{aligned} \quad (6)$$

where I is the length of the produced response; $g(\cdot)$ is an RNN; \mathbf{s}_i is the hidden state of the decoder at time step i which is conditioned on the previously generated token y_{i-1} , the previous hidden state \mathbf{s}_{i-1} , and the weighted attention vector




Figure 2: The training process of the generative model. First, the dialogue context is encoded: \mathbf{h} is the final hidden state of the context encoder. Then we derive the diversity-promoting latent variable \mathbf{z} . Next, we compute the latent variable c that corresponds to the measure of coherence between the dialogue context x and the generated response y . We concatenate all three vectors into \mathbf{s} to feed the decoder. \mathbf{a} is the attention matrix calculated for every time step of the decoding process.

\mathbf{a}_i :

$$\mathbf{s}_i = f(y_{i-1}, \mathbf{s}_{i-1}, \mathbf{a}_i) \quad (7)$$

$$\mathbf{a}_i = \sum_{j=1}^J \mathbf{w}_{ij} \mathbf{h}_i \quad (8)$$

where J is the number of tokens of the dialogue context; \mathbf{h}_i is the i^{th} hidden state of the encoder; the attention weight \mathbf{w}_{ij} of each context hidden state \mathbf{h}_i is computed following Luong et al. (2015).

Context Gate: To increase the influence of code c , we introduce the context gate k . Unlike Tu et al. (2017), whose context gate assigns an element-wise weight to the input signal deriving from the encoder RNN, we build the context gate conditioned only on the coherence signal:

$$k_i = \lambda \sigma(c - c'_i) \quad (9)$$

where σ is the sigmoid function; λ is a bias term;³ c is the target value of the measure of coherence, calculated by $C(\mathbf{x}, \mathbf{y})$ (see Section 2.1); c'_i is the measure of coherence between the dialogue context and the generated prefix sentence at time step i , calculated by $C(\mathbf{x}, \mathbf{y}^{<i})$. Now Eq (7) with the context gate applied to \mathbf{s}_i can be rewritten as:

$$\mathbf{s}_i = f(y_{i-1}, (1 - k_i) \circ \mathbf{s}_{i-1}, k_i \circ \mathbf{a}_i) \quad (10)$$

where \circ denotes element-wise multiplication.

The coherence-informed context gate aims to dynamically control the ratio at which preceding dialogue context and previously generated tokens of the current response contribute to the generation of the next token in the response.

³We set λ empirically against the development set.

2.3 Training

Our generator is trained similarly to Hu et al. (2017). The objective function is a weighted combination of three losses (generation, coherence, and diversity):

$$\mathcal{L} = \mathcal{L}_G + \lambda_c \mathcal{L}_c + \lambda_z \mathcal{L}_z \quad (11)$$

To teach the generator to produce responses close to the training data, we maximize the generation probability of the training response $\log p(\mathbf{y}|\mathbf{x})$ given the dialogue context according to Eq (2). During training, we set $\mathcal{L}_G = -\log p(\mathbf{y}|\mathbf{x})$ and minimize the following:

$$\mathcal{L}_G = D_{KL}(q(\mathbf{z}|\mathbf{x}, \mathbf{y}) \| p(\mathbf{z}|\mathbf{x})) - \mathbb{E}_{q(\mathbf{z}|\mathbf{x}, \mathbf{y})p(c|\mathbf{x}, \mathbf{y})} [\log p(\mathbf{y}|\mathbf{x}, \mathbf{z}, c)] \quad (12)$$

Apart from the generation loss, the coherence measure provides an extra learning signal \mathcal{L}_c which pushes the generator to produce responses that match the coherence signal given by the latent variable c .

$$\mathcal{L}_c = -\mathbb{E}_{p(\mathbf{z}|\mathbf{x})p(c)} \left[\log p(c|\mathbf{x}, \tilde{G}(\mathbf{x}, \mathbf{z}, c)) \right] \quad (13)$$

In Eq (13), $p(c) = N(0, 1)$ is the prior distribution of the coherence variable c . To ensure that the loss is differentiable, we cannot sample words from the response vocabulary. Instead we define $\tilde{G}(\mathbf{x}, \mathbf{z}, c) = \mathbf{y}^s = \{\mathbf{y}_1^s, \dots, \mathbf{y}_i^s, \dots, \mathbf{y}_I^s\}$ as the sequence of output word probability distributions. $p(c|\mathbf{x}, \tilde{G}(\mathbf{x}, \mathbf{z}, c))$ is predicted by the coherence measure defined in Eq (1) with \mathbf{y}^{emb} set as:

$$\mathbf{y}^{emb} = \sum_{i=1}^I \mathbf{y}_j^{s\top} M_{glv} \quad (14)$$

where M_{gloVe} is the word embedding matrix trained using GloVe (Section 2.1).

The last component in Eq (11) is the independent constraint \mathcal{L}_z that forces the soft distribution over the generated response \tilde{G} to be diverse, so that it is able to faithfully reproduce the latent variable \mathbf{z} :

$$\mathcal{L}_z = -\mathbb{E}_{p(\mathbf{z}|\mathbf{x})p(c)} \left[\log q \left(\mathbf{z}|\mathbf{x}, \tilde{G}(\mathbf{x}, \mathbf{z}, c) \right) \right] \quad (15)$$

where $q \left(\mathbf{z}|\mathbf{x}, \tilde{G}(\mathbf{x}, \mathbf{z}, c) \right)$ is predicted by the posterior network with \mathbf{y}_j^s as the soft input to the RNN encoder at each time step j .

2.4 Inference

Figure 3 shows the inference process of the generative model. Given a dialogue context \mathbf{x} and an expected coherence value c , the context encoder first encodes the dialogue context into a hidden state \mathbf{h} . The prior network then generates a sample \mathbf{z}' conditioned on the dialogue context. The decoder is initialized with \mathbf{s} , i.e., the concatenation of \mathbf{h} , \mathbf{z} and c . During decoding, the next word is generated via the context gate modulating between the attention-reweighted context and the previously generated words of the response.

3 Dataset and Filtering

Dataset for Generator

We train and evaluate our models on the OpenSubtitles corpus (Lison and Tiedemann, 2016) with automatic dialogue turn segmentation (Lison and Meena, 2016).⁴ A training pair consists of a dialogue context and a corresponding response. We consider three consecutive turns as the dialogue context and the following turn as the response. From a total of 65M instances, we select those that have context and response lengths of less than 120 and 30 words, respectively. We create two datasets:

1. **OST** (plain OpenSubtitles) consists of 2M/4K/4K instances as our training/development/test sets, selected randomly from the whole corpus;
2. **fOST** (filtered OpenSubtitles) contains the same amount of instances, but randomly selected only among those that have a measure of

⁴<http://www.opensubtitles.org>

Dataset	Coh	D-1%	D-2%	D-Sent%
OST	0.390	14.3	57.9	83.8
fOST	0.801	15.5	62.9	89.3

Table 1: Coherence and diversity metrics⁷ for the OST and fOST datasets (see Section 3 for the datasets and Section 4.2 for metrics definition).

coherence score $C(\mathbf{x}, \mathbf{y}) \geq 0.68$.⁵

Filtering of the OpenSubtitles corpus is motivated by the fact that by removing the video and audio modalities which the subtitles originally accompanied, we are very often left with incomplete and incoherent dialogues. Therefore, by keeping dialogues with high coherence scores, we aim at building a high quality corpus with (1) more semantically coherent and topically related contexts and responses, and (2) fewer general and dull responses. Table 3 shows the coherence and diversity metrics (cf. Section 4.2) between OST and fOST. Unsurprisingly, coherence for fOST is much higher than OST, with a slightly higher diversity. We list dialogue examples for different coherence scores in Supplemental Material B.

Dataset for Coherence Measure

In order to accurately measure coherence on our domain using the semantic distance as defined in Section 2.1, we train GloVe embeddings on the full OpenSubtitles corpus (i.e. 100K movies).

4 Experiments

Our generator model, ablative variants, and baselines are implemented using the publicly available OpenNMT-py framework (Klein et al., 2017) based on Bahdanau et al. (2015) and Luong et al. (2015). We used the publicly available glovepython package⁸ to implement our coherence measure.

We experiment on two versions of our model: (1) cVAE with the coherence context gate as described in Section 2.3 (cVAE-XGate), (2) cVAE with the original context gate implementation of

⁵The coherence score is calculated as shown in Eq (1). We observed that the scores on the training set follow a normal distribution with a slight tail on the negatively correlated side, so we fit a normal distribution to the data with parameters $N(0.25, 0.22)$ and set the cut-off to $+2\sigma$. A histogram of coherence scores is shown in Figure 5 in Supplemental Material A.

⁷Note that Distinct-1 and Distinct-2 are computed on a randomly selected subsets of 4k responses.

⁸<https://github.com/maciejkula/glove-python>




Figure 3: The inference process of the generative model, where the latent variable c is given as an input.

(Tu et al., 2017) (*cVAE-CGate*). For each of these, we consider the main variant where the input coherence measure c is preset to a fixed ideal value as estimated on development data (1.0 for OST and 0.95 for fOST), as well as an oracle variant where we use the true coherence measure between the context and the gold-standard response in the test set (indicated with “(C)” in Tables 2 and 3).

We compare against two baseline models: (1) a vanilla E-D with attention (*Attention*) (Luong et al., 2015); (2) an enhancement where output beams are rescored using the maximum mutual information anti-language model (MMI-antiLM) of Li et al. (2016a) (*MMI*).

4.1 Parameter Settings

We set our model parameters based on preliminary experiments on the development data.

We use 2-layer RNNs with LSTM cells (Hochreiter and Schmidhuber, 1997) with input/hidden dimension of 128 for both the context encoder and the decoder. The dropout rate is set to 0.2 and the Adam optimizer (Kingma and Ba, 2015) is used to update the parameters. A vocabulary of 25,000 words is shared between the encoder and the decoder.

Both the posterior network and prior network for the latent variable learning are built with 2-layer LSTM RNNs with input/hidden dimension of 64. The dimension of the latent variable z is set to 20. Same as for the encoder and decoder, the dropout rate is 0.2 and the Adam optimizer is used to update the parameters.

The window size for GloVe computation in our coherence measure is set to 10.

4.2 Evaluation metrics

We use a number of metrics to evaluate the outputs of our models:

- **BLEU, B1, B2, B3** – the word-overlap score against gold-standard responses (Papineni et al., 2002) used by the vast majority of recent

dialogue generation works (Zhao et al., 2017; Yao et al., 2017; Li et al., 2017a, 2016c; Sordani et al., 2015; Li et al., 2016a; Ghazvininejad et al., 2017). BLEU in this paper refers to the default BLEU-4, but we also report on lower n -gram scores (B1, B2, B3).⁹

- **Coh** – our novel GloVe-based coherence score calculated using Eq (1) showing the semantic distance of dialogue contexts and generated responses.
- **D-1, D-2, D-Sent** – common metrics used to evaluate the diversity of generated responses (e.g. Li et al., 2016a; Xu et al., 2017; Xing et al., 2017; Dhingra et al., 2017): the proportion of distinct unigrams, bigrams, and sentences in the outputs.

5 Results

All model variants described in Section 4 are trained on both OST and fOST datasets. Tables 2 and 3 present the scores of all models tested on the OST and fOST test sets, respectively. Note that in addition to testing the models on the respective test sections of their training datasets, we also test them on the other dataset (OST-trained models on fOST and vice-versa). This way, we can observe the performance of the fOST-trained models in more noisy contexts and see how good the OST-trained models are when evaluated against coherent responses only.

Given all the evaluated model variants, we can observe the effects and contributions of the individual components of our setup:

- **Data filtering:** The models trained on fOST consistently outperform the same models trained on OST – for all evaluation metrics and on both test sets. This shows that coherence-based training data filtering is generally beneficial.

⁹We use the Multi-BLEU script from OpenNMT to measure BLEU scores.

Training data	Model	BLEU%	B1%	B2%	B3%	Coh	D-1%	D-2%	D-Sent%
OST	Attention	1.32	10.92	3.85	2.10	0.293	3.4	14.2	25.6
	MMI	1.31	11.06	3.88	2.09	0.284	3.3	14.6	28.2
	cVAE-CGate (C)	1.58	11.86	4.45	2.48	0.311	4.1	15.0	28.2
	cVAE-XGate (C)	1.51	13.38	4.97	2.58	0.324	3.9	14.5	29.8
	cVAE-CGate (1.0)	1.60	17.08	5.78	2.86	0.404	5.0	27.1	79.7
	cVAE-XGate (1.0)	1.44	15.83	5.34	2.62	0.413	4.5	22.6	80.2
fOST	Attention	1.79	15.43	5.65	2.94	0.758	11.9	41.8	92.7
	MMI	1.99	16.24	6.06	3.22	0.764	11.9	44.5	95.8
	cVAE-CGate (C)	2.10	15.98	6.05	3.35	0.728	11.9	37.6	88.4
	cVAE-XGate (C)	1.85	16.44	5.94	3.07	0.706	10.3	31.2	80.4
	cVAE-CGate (0.95)	2.02	15.52	5.78	3.16	0.767	10.6	44.8	98.7
	cVAE-XGate (0.95)	1.64	14.43	5.20	2.70	0.745	9.0	36.9	98.7

Table 2: Evaluation results on the OST test set (see Section 4 for model description and Section 4.2 for metrics definition). Note that the cVAE-CGate(C) / cVAE-XGate(C) models use the true c value between the context and the gold response as input. Other cVAE-CGate / cVAE-XGate models use fixed values for c selected on dev sets shown in brackets. BLEU score reported here is BLEU-4; B1, B2 and B3 denote lower n -gram BLEU scores.

Training data	Model	BLEU%	B1%	B2%	B3%	Coh	D-1%	D-2%	D-Sent%
OST	Attention	0.86	8.34	2.79	1.45	0.284	3.6	14.6	29.4
	MMI	0.89	8.47	2.89	1.48	0.278	3.7	15.3	31.5
	cVAE-CGate (C)	1.64	10.20	4.17	2.40	0.329	5.1	19.4	35.8
	cVAE-XGate (C)	1.80	11.70	4.90	2.83	0.359	5.2	19.2	39.7
	cVAE-CGate (1.0)	2.25	16.82	6.81	3.70	0.422	5.4	28.2	81.0
	cVAE-XGate (1.0)	2.41	18.62	7.56	4.09	0.434	4.8	23.4	84.0
fOST	Attention	3.84	16.65	8.72	5.54	0.803	12.8	43.4	88.7
	MMI	3.84	16.81	8.78	5.57	0.803	12.6	42.5	88.8
	cVAE-CGate (C)	4.58	17.64	9.53	6.30	0.796	12.4	41.6	85.5
	cVAE-XGate (C)	4.33	18.43	9.59	6.11	0.783	10.7	33.1	78.8
	cVAE-CGate (0.95)	4.98	20.95	10.93	7.02	0.814	12.1	51.4	98.2
	cVAE-XGate (0.95)	4.47	20.98	10.43	6.50	0.797	10.4	42.5	97.6

Table 3: Evaluation results on the fOST test set (see Section 4 and Table 2 for model description; see Section 4.2 for metrics definition). BLEU score reported here is BLEU-4; B1, B2 and B3 denote lower n -gram BLEU scores.

- *cVAE-Context Gate models*: Nearly all cVAE-based models perform markedly better than the baselines w.r.t. BLEU, coherence, and diversity.¹⁰

If we look at models trained on OST and tested on fOST (the top half of Table 3), we can see that all cVAE-based models, especially cVAE-XGate, are able to learn to produce coherent and diverse response even when trained on a noisy, incoherent corpus. Examples of responses generated by the baseline MMI model and by cVAE-XGate in Figure 4 show that cVAE-XGate mostly produces more diverse and coherent responses than MMI.

- *Preset c vs. oracle models with gold-standard c* : Table 2 shows that on the noisy OST test set, cVAE-based models using the gold-standard

value of c achieve higher BLEU scores than models using preset c . This is expected since many gold-standard responses in the unfiltered set have a low coherence score; – the model can generate a more generic response if the gold-standard c is low. The models with preset c always attempt to generate coherent responses, which is apparent from the other metrics: Coh and D-Sent are consistently higher than for models using gold-standard c .

On the fOST test set where only high-coherence responses are expected, models using fixed c consistently reach higher scores in all metrics including BLEU (see Table 3). This shows that in general, using a preset constant value of c works well, even better than using the gold-standard c .

In sum, using our coherence measure both for data filtering and inside the models leads to output performance improvements.

¹⁰ We performed paired bootstrap re-sampling for the best cVAE model and the best baseline model in each experiments set (Table 2 and Table 3) as is done for MT (Koehn, 2004), which confirmed statistical significance at 99% confidence level for all cases except for models trained on fOST and tested on OST (bottom half of Table 2).

	Dialogue context	Response
1	A: i have an audition at 4:00. B: apparently now i'm <u>cooking dinner</u> . A: how did this all happen?	B-GT: you volunteered. B-MMI: i don't know. B-XGT: well , i'm going to <u>have a dinner</u> .
2	A: it's not that easy. B: we can't contact potential donors. A: if we want to <u>save frank</u> , his brother is his only hope	B-GT: the waiting list would take too long. B-MMI: i don't know. B-XGT: we can't <u>save him</u> .
3	A: that <u>fish</u> ate my <u>fish</u> ! B: you cost me £30,000 A: sorry, folks.	B-GT: you'll wait till next year to win one of these. B-MMI: i'm sorry. B-XGT: you owe me a little <u>fish</u> .
4	A: was it what you wanted? B: no. A: i wanted an <u>airplane</u> .	B-GT: you like model planes? B-MMI: i wanted to know. B-XGT: i wanted a <u>helicopter</u> .
5	A: great. B: thanks so much. A: <u>vanilla decaf latte</u> ?	B-GT: yeah, a vanilla decaf latte. B-MMI: no. B-XGT: <u>yeah</u> , that 's a <u>good coffee</u>
6	A: to make everybody happy. B: which means that sometimes A: you got to take a pile of poop	B-GT: and stick it in your mouth. B-MMI: and you got to take a shower. B-XGT: and you got to take a walk.
7	A: now you'll knock again and say you forgot to ask B: if she has a friend for your friend. A: are you <u>crazy</u> ?	B-GT: should i knock and make her come here for that? B-MMI: you're <u>crazy</u> ! B-XGT: i'm <u>crazy</u> , i'm <u>crazy</u> .
8	A: i knew you would say that. B: mr. kerem is a <u>charming and decent guy</u> . A: and he's really into you.	B-GT: you are all way too pretty to be arguing. B-MMI: he's a <u>very good man</u> . B-XGT: i didn't say that.

Figure 4: Generation output examples comparing two models trained on the OST dataset and tested on the fOST test set: the MMI baseline (B-MMI) and our cVAE-XGate (1.0) model (B-XGT). B-GT is the ground-truth response from the test set. The three sequential dialog turns on the left are the preceding dialogue context used to generate the responses. Corresponding topical phrases are underlined. We can see that cVAE-XGate (1.0) mostly produces markedly more coherent and specific outputs than MMI (1-5). In some cases, it is comparable with MMI (6-7) and occasionally, it is less coherent (8).

6 Related Work

Our work fits into the context of the very active area of end-to-end generative conversation models, where neural E-D approaches have been first applied by Vinyals and Le (2015) and extended by many others since.

Many works address the lack of diversity and coherence in E-D outputs (Sountsov and Sarawagi, 2016; Wei et al., 2017) but do not attempt to model coherence directly, unlike our work: Li et al. (2016a) use anti-LM reranking; Li et al. (2016c) modify the beam search decoding algorithm, similar to Shao et al. (2017) in addition to using a self-attention model. Mou et al. (2016) predict keywords for the output in a preprocessing step while Wu et al. (2018) preselect a vocabulary subset to be used for decoding. Li et al. (2016b) focus specifically on personality generation (using personality embeddings) and Wang et al. (2017) promote topic-specific outputs by language-model rescore and sampling.

A lot of recent works explore the use of additional training signals and VAE setups in dialogue

generation. In contrast to this paper, they do not focus explicitly on coherence: Asghar et al. (2017) use reinforcement learning with human-provided feedback, Li et al. (2017a) use a RL scenario with length as reward signal. Li et al. (2017b) add an adversarial discriminator to provide RL rewards (discriminating between human and machine outputs), Xu et al. (2017) use a full adversarial training setup. The most recent works explore the usage of VAEs: Cao and Clark (2017) explore a vanilla VAE setup conditioned on dual encoder (for contexts and responses) during training, the model of Serban et al. (2017) uses a VAE in a hierarchical E-D model. Shen et al. (2017) use a cVAE conditioned on sentiment and response genericity (based on a handwritten list of phrases). Shen et al. (2018) combine a cVAE with a plain VAE in an adversarial fashion.

We also draw on ideas from other areas than dialogue generation to build our models: Tu et al. (2017)'s context gates originate from machine translation and Hu et al. (2017)'s cVAE training stems from free-text generation.

7 Conclusions and Future Work

We showed that explicitly modeling coherence and optimizing towards coherence and diversity leads to better-quality outputs in dialogue response generation. We introduced three extensions to current encoder-decoder response generation models: (1) we defined a measure of coherence based on GloVe embeddings (Pennington et al., 2014), (2) we filtered the OpenSubtitles training corpus (Lison and Meena, 2016) based on this measure to obtain coherent and diverse training instances, (3) we trained a cVAE model based on (Hu et al., 2017) and (Tu et al., 2017) that uses our coherence measure as one of the training signals. Our experimental results showed a considerable improvement in the output quality over competitive models, which demonstrates the effectiveness of our approach.

In future work, we plan to replace the GloVe-based measure of coherence with a trained discriminator that distinguishes between coherent and incoherent responses (Li and Jurafsky, 2017). This will allow us to use extend the notion of coherence to account for phenomena such as topic shifts. We also plan to verify the results with a human evaluation study.

Acknowledgements

This research received funding from the EPSRC project MaDrIgAL (EP/N017536/1). The Titan Xp used for this research was donated by the NVIDIA Corporation.

References

- Nabiha Asghar, Pascal Poupart, Xin Jiang, and Hang Li. 2017. Deep Active Learning for Dialogue Generation. In *6th Joint Conference on Lexical and Computational Semantics (*SEM) 2017*. ArXiv: 1612.03929.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *3rd International Conference on Learning Representations (ICLR2015)*, San Diego, CA, USA. ArXiv: 1409.0473.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Kris Cao and Stephen Clark. 2017. Latent Variable Dialogue Models and their Diversity. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 182–187, Valencia, Spain. Association for Computational Linguistics.
- Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988.
- Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 484–495.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen tau Yih, and Michel Galley. 2017. A knowledge-grounded neural conversation model. *arXiv preprint arXiv:1702.01932*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *International Conference on Machine Learning*, pages 1587–1596.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations*, San Diego, CA, USA. ArXiv: 1412.6980.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. OpenNMT: Open-source toolkit for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, Vancouver, Canada.
- 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. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016b. A Persona-Based Neural

- Conversation Model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Berlin, Germany. ArXiv: 1603.06155.
- Jiwei Li and Dan Jurafsky. 2017. Neural Net Models for Open-Domain Discourse Coherence. In *Empirical Methods in Natural Language Processing*, Copenhagen, Denmark. ArXiv: 1606.01545.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016c. A simple, fast diverse decoding algorithm for neural generation. *arXiv preprint arXiv:1611.08562*.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2017a. Learning to decode for future success. *arXiv preprint arXiv:1701.06549*.
- Jiwei Li, Will Monroe, Tianlin Shi, Alan Ritter, and Dan Jurafsky. 2017b. Adversarial Learning for Neural Dialogue Generation. *arXiv:1701.06547 [cs]*. ArXiv: 1701.06547.
- Pierre Lison and Raveesh Meena. 2016. Automatic Turn Segmentation for Movie & TV Subtitles. In *2016 IEEE Workshop on Spoken Language Technology*. IEEE conference proceedings.
- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*, Portorož, Slovenia.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. ArXiv: 1508.04025.
- Lili Mou, Yiping Song, Rui Yan, Ge Li, Lu Zhang, and Zhi Jin. 2016. Sequence to Backward and Forward Sequences: A Content-Introducing Approach to Generative Short-Text Conversation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3349–3358, Osaka, Japan. The COLING 2016 Organizing Committee. ArXiv: 1607.00970.
- 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.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Iulian Serban, Alessandro Sordoni, Ryan Joseph Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural responding machine for short-text conversation. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Louis Shao, Stephan Gouws, Denny Britz, Anna Goldie, Brian Strope, and Ray Kurzweil. 2017. Generating Long and Diverse Responses with Neural Conversation Models. In *Empirical Methods in Natural Language Processing*, pages 2200–2009, Copenhagen, Denmark. ArXiv: 1701.03185.
- 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 (ACL 2017)*, Vancouver, Canada. ArXiv: 1705.00316.
- Xiaoyu Shen, Hui Su, Shuzi Niu, and Vera Demberg. 2018. Improving Variational Encoder-Decoders in Dialogue Generation. In *AAAI 2018*. ArXiv: 1802.02032.
- 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 *NAACL*.
- Pavel Sountsov and Sunita Sarawagi. 2016. Length bias in Encoder Decoder Models and a Case for Global Conditioning. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1516–1525, Austin, TX, USA. Association for Computational Linguistics.
- Zhaopeng Tu, Yang Liu, Zhengdong Lu, Xiaohua Liu, and Hang Li. 2017. Context Gates for Neural Machine Translation. *Transactions of the Association for Computational Linguistics*, 5:87–99.
- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. In *ICML Deep Learning Workshop*.
- 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, Copenhagen, Denmark.
- Bolin Wei, Shuai Lu, Lili Mou, Hao Zhou, Pascal Poupart, Ge Li, and Zhi Jin. 2017. Why Do Neural Dialog Systems Generate Short and Meaningless Replies? A Comparison between Dialog and Translation. *arXiv:1712.02250 [cs]*. ArXiv: 1712.02250.

- Yu Wu, Wei Wu, Dejian Yang, Can Xu, Zhoujun Li, and Ming Zhou. 2018. Neural Response Generation with Dynamic Vocabularies. In *AAAI 2018*, New Orleans, LA, USA. ArXiv: 1711.11191.
- 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.
- Zhen Xu, Bingquan Liu, Baoxun Wang, SUN Chengjie, Xiaolong Wang, Zhuoran Wang, and Chao Qi. 2017. Neural response generation via gan with an approximate embedding layer. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 628–637.
- Lili Yao, Yaoyuan Zhang, Yansong Feng, Dongyan Zhao, and Rui Yan. 2017. Towards implicit content-introducing for generative short-text conversation systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2180–2189.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders. In *ACL*, Vancouver, Canada. ArXiv: 1703.10960.