

# Variational Hierarchical Dialog Autoencoder for Dialog State Tracking Data Augmentation

Kang Min Yoo<sup>1</sup> Hanbit Lee<sup>1</sup> Franck Dernoncourt<sup>2</sup> Trung Bui<sup>2</sup>  
Walter Chang<sup>2</sup> Sang-goo Lee<sup>1</sup>

<sup>1</sup>Seoul National University, Seoul, Korea

<sup>2</sup>Adobe Research, San Jose, CA, USA

{kangminyoo, skcheon, sglee}@europa.snu.ac.kr

{dernonco, bui, wachang}@adobe.com

## Abstract

Recent works have shown that generative data augmentation, where synthetic samples generated from deep generative models complement the training dataset, benefit NLP tasks. In this work, we extend this approach to the task of dialog state tracking for goal-oriented dialogs. Due to the inherent hierarchical structure of goal-oriented dialogs over utterances and related annotations, the deep generative model must be capable of capturing the coherence among different hierarchies and types of dialog features. We propose the Variational Hierarchical Dialog Autoencoder (VHDA) for modeling the complete aspects of goal-oriented dialogs, including linguistic features and underlying structured annotations, namely speaker information, dialog acts, and goals. The proposed architecture is designed to model each aspect of goal-oriented dialogs using inter-connected latent variables and learns to generate coherent goal-oriented dialogs from the latent spaces. To overcome training issues that arise from training complex variational models, we propose appropriate training strategies. Experiments on various dialog datasets show that our model improves the downstream dialog trackers' robustness via generative data augmentation. We also discover additional benefits of our unified approach to modeling goal-oriented dialogs – dialog response generation and user simulation, where our model outperforms previous strong baselines.

## 1 Introduction

Data augmentation, a technique that augments the training set with label-preserving synthetic samples, is commonly employed in modern machine learning approaches. It has been used extensively in visual learning pipelines (Shorten and Khoshgof-taar, 2019) but less frequently for NLP tasks due to the lack of well-established techniques in the

area. While some notable work exists in text classification (Zhang et al., 2015), spoken language understanding (Yoo et al., 2019), and machine translation (Fadaee et al., 2017), we still lack the full understanding of utilizing generative models for text augmentation.

Ideally, a data augmentation technique for supervised tasks must synthesize *distribution-preserving* and *sufficiently realistic* samples. Current approaches for data augmentation in NLP tasks mostly revolve around thesaurus data augmentation (Zhang et al., 2015), in which words that belong to the same semantic role are substituted with one another using a preconstructed lexicon, and noisy data augmentation (Wei and Zou, 2019) where random editing operations create perturbations in the language space. Thesaurus data augmentation requires a set of handcrafted semantic dictionaries, which are costly to build and maintain, whereas noisy data augmentation does not synthesize sufficiently realistic samples. The recent trend (Hu et al., 2017; Yoo et al., 2019; Shin et al., 2019) gravitates towards *generative data augmentation* (GDA), a class of techniques that leverage deep generative models such as VAEs to delegate the automatic discovery of novel class-preserving samples to machine learning. In this work, we explore GDA in the context of dialog modeling and contextual understanding.

Goal-oriented dialogs occur between a user and a system that communicates verbally to accomplish the user's goals (Table 6). However, because the user's goals and the system's possible actions are not transparent to each other, both parties must rely on verbal communications to infer and take appropriate actions to resolve the goals. Dialog state tracker is a core component of such systems, enabling it to track the dialog's latest status (Henderson et al., 2014). A dialog state typically consists of `inform` and `request` types of slot values.

For example, a user may verbally refer to a previously mentioned food type as the preferred one - e.g., Asian (`inform(food=asian)`). Given the user utterance and historical turns, the state tracker must infer the user’s current goals. As such, we can view dialog state tracking as a sparse sequential multi-class classification problem. Modeling goal-oriented dialogs for GDA requires a novel approach that simultaneously solves state tracking, user simulation (Schatzmann et al., 2007), and utterance generation.

Various deep models exist for modeling dialogs. The Markov approach (Serban et al., 2017) employs a sequence-to-sequence variational autoencoder (VAE) (Kingma and Welling, 2013) structure to predict the next utterance given a deterministic context representation, while the holistic approach (Park et al., 2018) utilizes a set of global latent variables to encode the entire dialog, improving the awareness in general dialog structures. However, current approaches are limited to linguistic features. Recently, Bak and Oh (2019) proposed a hierarchical VAE structure that incorporates the speaker’s information, but we have yet to explore a universal approach for encompassing fundamental aspects of goal-oriented dialogs. Such a unified model capable of disentangling latents into specific dialog aspects can increase the modeling efficiency and enable interesting extensions based on the fine-grained controllability.

This paper proposes a novel multi-level hierarchical and recurrent VAE structure called Variational Hierarchical Dialog Autoencoder (VHDA). Our model enables modeling all aspects (speaker information, goals, dialog acts, utterances, and general dialog flow) of goal-oriented dialogs in a disentangled manner by assigning latents to each aspect. However, complex and autoregressive VAEs are known to suffer from the risk of *inference collapse* (Cremer et al., 2018), in which the model converges to a local optimum where the generator network neglects the latents, reducing the generation controllability. To mitigate the issue, we devise two simple but effective training strategies.

Our contributions are summarized as follows.

1. We propose a novel deep latent model for modeling dialog utterances and their relationships with the goal-oriented annotations. We show that the strong level of coherence and accuracy displayed by the model allows it to be used for augmenting dialog state tracking

datasets.

2. Leveraging the model’s generation capabilities, we show that generative data augmentation is attainable even for the complex dialog-related tasks that pertain to both hierarchical and sequential annotations.
3. We propose simple but effective training policies for our VAE-based model, which have applications in other similar VAE structures.

The code for reproducing this paper is available at [github](https://github.com/kaniblu/vhda)<sup>1</sup>.

## 2 Background and Related Work

**Dialog State Tracking.** Dialog state tracking (DST) predicts the user’s current goals and dialog acts, given the dialog context. Historically, DST models have gradually evolved from hand-crafted finite-state automata and multi-stage models (Dybkjær and Minker, 2008; Thomson and Young, 2010; Wang and Lemon, 2013) to end-to-end models that directly predict dialog states from dialog features (Zilka and Jurcicek, 2015; Mrkšić et al., 2017; Zhong et al., 2018; Nouri and Hosseini-Asl, 2018).

Among the proposed models, Neural Belief Tracker (NBT) (Mrkšić et al., 2017) decreases reliance on handcrafted semantic dictionaries by reformulating the classification problem. Globally Self-attentive Dialog tracker (GLAD) (Zhong et al., 2018) introduces global modules for sharing parameters across slots and local modules, allowing the learning of slot-specific feature representations. Globally-Conditioned Encoder (GCE) (Nouri and Hosseini-Asl, 2018) improves further by forgoing the separation of global and local modules, allowing the unified module to take slot embeddings for distinction. Recently, dialog state trackers based on pre-trained language models have demonstrated their strong performance in many DST tasks (Wu et al., 2019; Kim et al., 2019; Hosseini-Asl et al., 2020). While the utilization of large-scale pre-trained language models is not within our scope, we wish to explore further concerning the recent advances in the area.

**Conversation Modeling.** While the previous approaches for hierarchical dialog modeling relate to the Markov assumption (Serban et al., 2017), recent approaches have geared towards utilizing

<sup>1</sup><https://github.com/kaniblu/vhda>

global latent variables for representing the holistic dialog structure (Park et al., 2018; Gu et al., 2018; Bak and Oh, 2019), which helps in preserving long-term dependencies and total semantics. In this work, we employ global latent variables to maximize the effectiveness in preserving dialog semantics for data augmentation.

**Data Augmentation.** Transformation-based data augmentation is popular in vision learning (Shorten and Khoshgoftaar, 2019) and speech signal processing (Ko et al., 2015), while thesaurus and noisy data augmentation techniques are more common for text. (Zhang et al., 2015; Wei and Zou, 2019). Recently, generative data augmentation (GDA), augmenting data gather from samples generated from fine-tuned deep generative models, have gained traction in several NLP tasks (Hu et al., 2017; Hou et al., 2018; Yoo et al., 2019; Shin et al., 2019). GDA can be seen as a form of unsupervised data augmentation, delegating the automatic discovery of novel data to machine learning without injecting external knowledge or data sources. While most works utilize VAE for the generative model, some works achieved a similar effect without employing variational inference (Kurata et al., 2016; Hou et al., 2018). In contrast to unsupervised data augmentation, another line of work has explored self-supervision mechanisms to fine-tune the generators for specific tasks (Tran et al., 2017; Antoniou et al., 2017; Cubuk et al., 2018). Recent work proposed a reinforcement learning-based noisy data augmentation framework for state tracking (Yin et al., 2019). Our work belongs to the family of unsupervised GDA, which can incorporate self-supervision mechanisms. We wish to explore further in this regard.

### 3 Proposed Model

This section describes VHDA, our latent variable model for generating goal-oriented dialog datasets. We first introduce a set of notations for describing core concepts.

#### 3.1 Notations

A dialog dataset  $\mathbb{D}$  is a set of  $N$  i.i.d samples  $\{\mathbf{c}_1, \dots, \mathbf{c}_N\}$ , where each  $\mathbf{c}$  is a sequence of turns  $(\mathbf{v}_1, \dots, \mathbf{v}_T)$ . Each goal-oriented dialog turn  $\mathbf{v}$  is a tuple of speaker information  $\mathbf{r}$ , the speaker’s goals  $\mathbf{g}$ , dialog state  $\mathbf{s}$ , and the speaker’s utterance  $\mathbf{u}$ :  $\mathbf{v} = (\mathbf{r}, \mathbf{g}, \mathbf{s}, \mathbf{u})$ . Each utterance  $\mathbf{u}$  is a sequence of words  $(w_1, \dots, w_{|\mathbf{u}|})$ . Goals  $\mathbf{g}$  or




Figure 1: Graphical representation of VHDA. Solid and dashed arrows represent generation and recognition respectively.

a dialog state  $\mathbf{s}$  is defined as a set of the smallest unit of dialog act specification  $a$  (Henderson et al., 2014), which is a tuple of dialog act, slot and value defined over the space of  $\mathcal{T}$ ,  $\mathcal{S}$ , and  $\mathcal{V}$ :  $\mathbf{g} = \{a_1, \dots, a_{|\mathbf{g}|}\}$ ,  $\mathbf{s} = \{a_1, \dots, a_{|\mathbf{s}|}\}$ , where  $a_i \in \mathcal{A} = (\mathcal{T}, \mathcal{S}, \mathcal{V})$ . A dialog act specification is represented as  $\langle \text{act} \rangle (\langle \text{slot} \rangle = \langle \text{value} \rangle)$ .

#### 3.2 VHCR

Given a conversation  $\mathbf{c}$ , Variational Hierarchical Conversational RNN (VHCR) (Park et al., 2018) models the holistic features of the conversation and the individual utterances  $\mathbf{u}$  using a hierarchical and recurrent VAE model. The model introduces global-level latent variables  $\mathbf{z}^{(c)}$  for encoding the high-level dialog structure and, at each turn  $t$ , local-level latent variables  $\mathbf{z}_t^{(u)}$  responsible for encoding and generating the utterance at turn  $t$ . The local latent variables  $\mathbf{z}^{(u)}$  conditionally depends on  $\mathbf{z}^{(c)}$  and previous observations, forming a hierarchical structure with the global latents. Furthermore, hidden variables  $\mathbf{h}_t$ , which are conditionally dependent on the global information and the hidden variables from the previous step  $\mathbf{h}_{t-1}$ , facilitate the latent inference.

#### 3.3 VHDA

We propose Variational Hierarchical Dialog Autoencoder (VHDA) to generate dialogs and their underlying dialog annotations simultaneously (Figure 1). Like VHCR, we employ a hierarchical VAE structure to capture holistic dialog semantics using the conversation latent variables  $\mathbf{z}^{(c)}$ . Our model

incorporates full dialog features using turn-level latents  $\mathbf{z}^{(r)}$  (speaker),  $\mathbf{z}^{(g)}$  (goal),  $\mathbf{z}^{(s)}$  (dialog state), and  $\mathbf{z}^{(u)}$  (utterance), motivated by speech act theory (Searle et al., 1980). Specifically, at a given dialog turn, the information about the speaker, the speaker’s goals, the speaker’s turn-level dialog acts, and the utterance all cumulatively determine one after the other in that order.

VHDA consists of multiple encoder and decoder modules, each responsible for encoding or generating a particular dialog feature. The encoders share the identical sequence-encoding architecture described as follows.

**Sequence Encoder Architecture.** Given a sequence of variable number of elements  $\mathbf{X} = [\mathbf{x}_1; \dots; \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$ , where  $n$  is the number of elements, the goal of a sequence encoder is to extract a fixed-size representation  $\mathbf{h} \in \mathbb{R}^d$ , where  $d$  is the dimensionality of the hidden representation. For our implementation, we employ the self-attention mechanism over hidden outputs of bidirectional LSTM (Hochreiter and Schmidhuber, 1997) cells produced from the input sequence. We also allow the attention mechanism to be optionally queried by  $\mathbf{Q}$ , enabling the sequence to depend on external conditions, such as using the dialog context to attend over an utterance:

$$\begin{aligned} \mathbf{H} &= [\overrightarrow{\text{LSTM}}(\mathbf{X}); \overleftarrow{\text{LSTM}}(\mathbf{X})] \in \mathbb{R}^{n \times d} \\ \mathbf{a} &= \text{softmax}([\mathbf{H}; \mathbf{Q}]\mathbf{w} + b) \in \mathbb{R}^n \\ \mathbf{h} &= \mathbf{H}^\top \mathbf{a} \in \mathbb{R}^d. \end{aligned}$$

Here,  $\mathbf{Q} \in \mathbb{R}^{n \times d_q}$  is a collection of query vectors of size  $d_q$  where each vector corresponds to one element in the sequence;  $\mathbf{w} \in \mathbb{R}^{d+d_q}$  and  $b \in \mathbb{R}$  are learnable parameters. We encapsulate the above operations with the following notation:

$$\mathcal{E} : \mathbb{R}^{n \times d} (\times \mathbb{R}^{n \times d_q}) \rightarrow \mathbb{R}^d.$$

Our model utilizes the  $\mathcal{E}$  structure for encoding dialog features of variable lengths.

**Encoder Networks.** Based on the  $\mathcal{E}$  architecture, feature encoders are responsible for encoding dialog features from their respective raw feature spaces to hidden representations. For goals and turn states, the encoding consists of two steps. Initially, the multi-purpose dialog act encoder  $\mathcal{E}^{(a)}$  processes each dialog act triple of the goals  $a^{(g)} \in \mathbf{g}$  and turn states  $a^{(s)} \in \mathbf{s}$  into a fixed-size representation  $\mathbf{h}^{(a)} \in \mathbb{R}^{d^{(a)}}$ . The encoder treats the dialog act triples as sequences of tokens. Subsequently,

the goal encoder and the turn state encoder process those dialog act representations to produce goal representations and turn state representations, respectively:

$$\begin{aligned} \mathbf{h}^{(g)} &= \mathcal{E}^{(g)}([\mathcal{E}^{(a)}(a_1^{(g)}); \dots; \mathcal{E}^{(a)}(a_{|\mathbf{g}|}^{(g)})]) \\ \mathbf{h}^{(s)} &= \mathcal{E}^{(s)}([\mathcal{E}^{(a)}(a_1^{(s)}); \dots; \mathcal{E}^{(a)}(a_{|\mathbf{s}|}^{(s)})]). \end{aligned}$$

Note that, as the model is sensitive to the order of the dialog acts, we randomize the order during training to prevent overfitting. The utterances are encoded using the utterance encoder from the word embeddings space:  $\mathbf{h}^{(u)} = \mathcal{E}^{(u)}([\mathbf{w}_1; \dots; \mathbf{w}_{|u|}])$ , while the entire conversation is encoded by the conversation encoder from the encoded utterance vectors:  $\mathbf{h}^{(c)} = \mathcal{E}^{(c)}([\mathbf{h}_1^{(u)}; \dots; \mathbf{h}_T^{(u)}])$ . All sequence encoders mentioned above depend on the global latent variables  $\mathbf{z}^{(c)}$  via the query vector. For the speaker information, we use the speaker embedding matrix  $\mathbf{W}^{(r)} \in \mathbb{R}^{n^{(r)} \times d^{(r)}}$  to encode the speaker vectors  $\mathbf{h}^{(r)}$ , where  $n^{(r)}$  is the number of participants and  $d^{(r)}$  is the embedding size.

**Main Architecture.** At the top level, our architecture consists of five  $\mathcal{E}$  encoders, a context encoder  $\mathcal{C}$ , and four types of decoder  $\mathcal{D}$ . The context encoder  $\mathcal{C}$  is different from the other encoders, as it does *not* utilize the bidirectional  $\mathcal{E}$  architecture but a uni-directional LSTM cell. The four decoders  $\mathcal{D}^{(r)}$ ,  $\mathcal{D}^{(g)}$ ,  $\mathcal{D}^{(s)}$ , and  $\mathcal{D}^{(u)}$  generate respective dialog features.

$\mathcal{C}$  is responsible for keeping track of the dialog context by encoding all features generated so far. The context vector at  $t$  ( $\mathbf{h}_t$ ) is updated using the historical information from the previous step:

$$\begin{aligned} \mathbf{v}_{t-1} &= [\mathbf{h}_{t-1}^{(r)}; \mathbf{h}_{t-1}^{(g)}; \mathbf{h}_{t-1}^{(s)}; \mathbf{h}_{t-1}^{(u)}] \\ \mathbf{h}_t &= \mathcal{C}(\mathbf{h}_{t-1}, \mathbf{v}_{t-1}) \end{aligned}$$

where  $\mathbf{v}_t$  is represents all features at the step  $t$ .

VHDA uses the context information to successively generate turn-level latent variables using a series of generator networks:

$$\begin{aligned} p_\theta(\mathbf{z}_t^{(r)} | \mathbf{h}_t, \mathbf{z}^{(c)}) &= \mathcal{N}(\mu_t^{(r)}, \sigma_t^{(r)} \mathbf{I}) \\ p_\theta(\mathbf{z}_t^{(g)} | \mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{z}_t^{(r)}) &= \mathcal{N}(\mu_t^{(g)}, \sigma_t^{(g)} \mathbf{I}) \\ p_\theta(\mathbf{z}_t^{(s)} | \mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{z}_t^{(r)}, \mathbf{z}_t^{(g)}) &= \mathcal{N}(\mu_t^{(s)}, \sigma_t^{(s)} \mathbf{I}) \\ p_\theta(\mathbf{z}_t^{(u)} | \mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{z}_t^{(r)}, \mathbf{z}_t^{(g)}, \mathbf{z}_t^{(s)}) &= \mathcal{N}(\mu_t^{(u)}, \sigma_t^{(u)} \mathbf{I}) \end{aligned}$$

where all latents are assumed to be Gaussian. In addition, we assume the standard Gaussian for



the global latents:  $p(\mathbf{z}^{(c)}) = \mathcal{N}(0, \mathbf{I})$ . We implemented the Gaussian distribution encoders ( $\mu$  and  $\sigma$ ) using fully-connected networks  $f$ . We also apply softplus on the output of the networks to infer the variance of the distributions. Employing the reparameterization trick (Kingma and Welling, 2013) allows standard backpropagation during training of our model.

**Approximate Posterior Networks.** We use a separate set of parameters  $\phi$  and encoders to approximate the posterior distributions of latent variables from the evidence. In particular, the model infers the global latents  $\mathbf{z}^{(c)}$  using the conversation encoder  $\mathcal{E}^{(c)}$  solely from the linguistic features:

$$q_\phi(\mathbf{z}^{(c)} | \mathbf{h}_1^{(u)}, \dots, \mathbf{h}_T^{(u)}) = \mathcal{N}(\mu^{(c)}, \sigma^{(c)} \mathbf{I}).$$

Similarly, the approximate posterior distributions of all turn-level latent variables are estimated from the evidence in cascade, while maintaining the global conditioning:

$$\begin{aligned} q_\phi(\mathbf{z}_t^{(r)} | \mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{h}_t^{(r)}) &= \mathcal{N}(\mu_t^{(r)}, \sigma_t^{(r)} \mathbf{I}) \\ q_\phi(\mathbf{z}_t^{(g)} | \mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{z}_t^{(r)}, \mathbf{h}_t^{(g)}) &= \mathcal{N}(\mu_t^{(g)}, \sigma_t^{(g)} \mathbf{I}) \\ q_\phi(\mathbf{z}_t^{(s)} | \mathbf{h}_t, \dots, \mathbf{z}_t^{(g)}, \mathbf{h}_t^{(s)}) &= \mathcal{N}(\mu_t^{(s)}, \sigma_t^{(s)} \mathbf{I}) \\ q_\phi(\mathbf{z}_t^{(u)} | \mathbf{h}_t, \dots, \mathbf{z}_t^{(s)}, \mathbf{h}_t^{(u)}) &= \mathcal{N}(\mu_t^{(u)}, \sigma_t^{(u)} \mathbf{I}), \end{aligned}$$

where all Gaussian parameters are estimated using fully-connected layers, parameterized by  $\phi$ .

**Realization Networks.** A series of generator networks successively decodes dialog features from their respective latent spaces to realize the surface forms:

$$\begin{aligned} p_\theta(\mathbf{r}_t | \mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{z}_t^{(r)}) &= \mathcal{D}_\theta^{(r)}(\mathbf{h}_t, \mathbf{z}^{(c)}, \mathbf{z}_t^{(r)}) \\ p_\theta(\mathbf{g}_t | \mathbf{h}_t, \dots, \mathbf{z}_t^{(g)}) &= \mathcal{D}_\theta^{(g)}(\mathbf{h}_t, \dots, \mathbf{z}_t^{(g)}) \\ p_\theta(\mathbf{s}_t | \mathbf{h}_t, \dots, \mathbf{z}_t^{(s)}) &= \mathcal{D}_\theta^{(s)}(\mathbf{h}_t, \dots, \mathbf{z}_t^{(s)}) \\ p_\theta(\mathbf{u}_t | \mathbf{h}_t, \dots, \mathbf{z}_t^{(u)}) &= \mathcal{D}_\theta^{(u)}(\mathbf{h}_t, \dots, \mathbf{z}_t^{(u)}). \end{aligned}$$

The utterance decoder  $\mathcal{D}^{(u)}$  is implemented using the LSTM cell. To alleviate sparseness in goals and turn-level dialog acts, we formulate the classification problem as a set of binary classification problems (Mrkšić et al., 2017). Specifically, given a candidate dialog act  $a$ ,

$$p_\theta(a \in \mathbf{s}_t | \mathbf{v}_{<t}, \dots) = \sigma(\mathbf{o}_t^{(s)} \cdot \mathcal{E}^{(a)}(a))$$

where  $\sigma$  is the sigmoid function and  $\mathbf{o}_t^{(s)} \in \mathbb{R}^{d(a)}$  is the output of a feedforward network parameterized by  $\theta$  that predicts the dialog act specification embeddings. Goals are predicted analogously.

### 3.4 Training Objective

Given all the latent variables  $\mathbf{z}$  in our model, we optimize the evidence lower-bound (ELBO) of the goal-oriented dialog samples  $\mathbf{c}$ :

$$\begin{aligned} \mathcal{L}_{\text{VHDA}} &= \mathbb{E}_{q_\phi}[\log p_\theta(\mathbf{c} | \mathbf{z})] \\ &\quad - D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{c}) \| p(\mathbf{z})). \end{aligned} \quad (1)$$

The reconstruction term of Equation 5 can be factorized into posterior probabilities in the realization networks. Similarly, the KL-divergence term can be factorized and reformulated in approximate posterior networks and conditional priors based on the graphical structure.

### 3.5 Minimizing Inference Collapse

Inference collapse is a relatively common phenomenon among autoregressive VAE structures (Zhao et al., 2017). The hierarchical and recurrent nature of our model makes it especially vulnerable. The standard treatment for alleviating the inference collapse problem include (1) annealing the KL-divergence term weight during the initial training stage and (2) employing word dropouts on the decoder inputs (Bowman et al., 2016). For our model, we observe that the basic techniques are insufficient (Table 3). While more recent treatments exist (Kim et al., 2018; He et al., 2019), they incur high computational costs that prohibit practical deployment in our cases. We introduce two simpler but effective methods to prevent encoder degeneration.

**Mutual Information Maximization.** The KL-divergence term in the standard VAE ELBO can be decomposed to reveal the mutual information term (Hoffman and Johnson, 2016):

$$\begin{aligned} \mathbb{E}_{p_d}[D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{x}) \| p(\mathbf{z}))] &= \\ D_{\text{KL}}(q_\phi(\mathbf{z}) \| p(\mathbf{z})) &+ I_{q_\phi}(\mathbf{x}; \mathbf{z}) \end{aligned}$$

where  $p_d$  is the empirical distribution of the data. Re-weighting the decomposed terms for optimizing the VAE behaviors has been explored previously (Chen et al., 2018; Zhao et al., 2017; Tolstikhin et al., 2018). In this work, we propose simply canceling out the mutual information term by performing mutual information estimation as a post-procedure. Since the preservation of the conversation encoder  $\mathcal{E}^{(c)}$  and global latents is vital for generation controllability, we specifically maximize mutual information between the global latents and

the evidence:

$$\begin{aligned} \mathcal{L}_{\text{VHDA}} = & \mathbb{E}_{q_\phi} [\log p_\theta(\mathbf{c} \mid \mathbf{z})] \\ & - D_{\text{KL}}(q_\phi(\mathbf{z} \mid \mathbf{c}) \parallel p(\mathbf{z})) + I_{q_\phi}(\mathbf{c}; \mathbf{z}^{(c)}). \end{aligned} \quad (2)$$

In our work, the mutual information term is computed empirically using the Monte-Carlo estimator for each mini-batch. The details are provided in the supplementary material.

**Hierarchically-scaled Dropout.** Extending word dropouts and utterance dropouts Park et al. (2018), we apply dropouts discriminatively to all dialog features (goals and dialog acts) according to the feature hierarchy level. We hypothesize that employing dropouts could be detrimental to the learning of lower-level latent variables, as information dropouts stack multiplicatively along the hierarchy. However, it is also necessary in order to encourage meaningful encoding of latent variables. Specifically, we propose a novel dropout scheme that scales exponentially along with the hierarchical depth, allowing higher-level information to flow towards lower levels easily. For our implementation, we set the dropout ratio between two adjacent levels to 1.5, resulting in the dropout probabilities of [0.1, 0.15, 0.23, 0.34, 0.51] for speaker information to utterances. We confirm our hypothesis in § 4.2.

## 4 Experiments

### 4.1 Experimental Settings

Following the protocol in (Yoo et al., 2019), we generate three independent sets of synthetic dialog samples, and, for each augmented dataset, we repeatedly train the same dialog state tracker three times with different seeds. We compare the aggregated results from all nine trials with the baseline results. Ultimately, we repeat this procedure for all combinations of state trackers and datasets. For non-augmented baselines, we repeat the experiments ten times.

**Implementation Details.** The hidden size of dialog vectors is 1000, and the hidden size of utterance, dialog act specification, turn state, and turn goal representations is 500. The dimensionality for latent variables is between 100 and 200. We use GloVe (Pennington et al., 2014) and character (Hashimoto et al., 2017) embeddings as pre-trained word embeddings (400 dimensions total) for word and dialog act tokens. All models used Adam optimizer (Kingma and Ba, 2014) with the initial learning rate of  $1e-3$ . We annealed the KL-divergence

weights over 250,000 training steps. For data synthesis, we employ ancestral sampling to generate samples from the empirical posterior distribution. We fixed the ratio of synthetic to original data samples to 1.

**Datasets.** We conduct experiments on four state tracking corpora: *WoZ2.0* (Wen et al., 2017), *DSTC2* (Henderson et al., 2014), *MultiWoZ* (Budzianowski et al., 2018), and *DialEdit* (Manuvinakurike et al., 2018). These corpora cover a variety of domains (restaurant booking, hotel reservation, and image editing). Note that, because the MultiWoZ dataset is a multi-domain corpus, we extract single-domain dialog samples from the two most prominent domains (hotel and restaurant, denoted by *MultiWoZ-H* and *MultiWoZ-R*, respectively).

**Dialog State Trackers.** We use GLAD and GCE as the two competitive baselines for state tracking. Besides, modifications are applied to these trackers to stabilize the performance on random seeds (denoted as GLAD<sup>+</sup> and GCE<sup>+</sup>). Specifically, we enrich the word embeddings with subword information (Bojanowski et al., 2017) and apply dropout on word embeddings (dropout rate of 0.2). Furthermore, we also conduct experiments on a simpler architecture that shares a similar structure with GCE but does not employ self-attention for the sequence encoders (denoted as RNN).

**Evaluation Measures.** *Joint goal accuracy* (goal for short) measures the ratio of the number of turns whose goals a tracker has correctly identified over the total number of turns. Similarly, *request accuracy*, or *request*, measures the turn-level accuracy of request-type dialog acts, while *inform accuracy* (*inform*) measures the turn-level accuracy of inform-type dialog acts. Turn-level goals accumulate from inform-type dialog acts starting from the beginning of the dialog until respective dialog turns, and thus they can be inferred from historical inform-type dialog acts (Table 6).

### 4.2 Data Augmentation Results

**Main Results.** We present the data augmentation results in Table 1. The results strongly suggest that generative data augmentation for dialog state tracking is a viable strategy for improving existing DST models without modifying them, as improvements were observed at statistically significant levels regardless of the tracker and dataset.

The margin of improvements was more signifi-

GDA	MODEL	WoZ2.0		DSTC2		MWOZ-R		MWOZ-H		DIALEDIT	
		GOAL	REQ	GOAL	REQ	GOAL	INF	GOAL	INF	GOAL	REQ
-	RNN	74.5	96.1	69.7	96.0	43.7	69.4	25.7	55.6	35.8	96.6
VHDA	RNN	<b>78.7</b> <sup>‡</sup>	<b>96.7</b> <sup>‡</sup>	<b>74.2</b> <sup>†</sup>	<b>97.0</b> <sup>‡</sup>	<b>49.6</b> <sup>†</sup>	<b>73.4</b> <sup>†</sup>	<b>31.0</b> <sup>†</sup>	<b>59.7</b> <sup>†</sup>	<b>36.4</b> <sup>†</sup>	<b>96.8</b>
-	GLAD <sup>+</sup>	87.8	<b>96.8</b>	74.5	96.4	58.9	76.3	33.4	58.9	35.9	96.7
VHDA	GLAD <sup>+</sup>	<b>88.4</b>	96.6	<b>75.5</b> <sup>‡</sup>	<b>96.8</b> <sup>†</sup>	<b>61.5</b> <sup>†</sup>	<b>77.4</b>	<b>37.8</b> <sup>‡</sup>	<b>61.3</b> <sup>‡</sup>	<b>37.1</b> <sup>†</sup>	<b>96.8</b>
-	GCE <sup>+</sup>	88.7	97.0	74.8	96.3	60.5	76.7	36.5	61.0	36.1	96.6
VHDA	GCE <sup>+</sup>	<b>89.3</b> <sup>‡</sup>	<b>97.1</b>	<b>76.0</b> <sup>‡</sup>	<b>96.7</b> <sup>†</sup>	<b>63.3</b>	<b>77.2</b>	<b>38.3</b>	<b>63.1</b> <sup>†</sup>	<b>37.6</b> <sup>†</sup>	<b>96.8</b>

<sup>†</sup>  $p < 0.1$     <sup>‡</sup>  $p < 0.01$

Table 1: Results of data augmentation using VHDA for dialog state tracking on various datasets and state trackers. Note that we report inform accuracies for MultiWoZ datasets instead, as request-type prediction is trivial for those.

GOAL	DST	WoZ2.0		DSTC2	
		GOAL	REQ	GOAL	REQ
w/o	RNN	77.8	96.4	71.2	<b>97.2</b>
w/	RNN	<b>78.7</b>	<b>96.7</b>	<b>74.2</b>	97.0
w/o	GLAD <sup>+</sup>	86.5	<b>96.9</b>	74.7	<b>97.0</b>
w/	GLAD <sup>+</sup>	<b>88.4</b>	96.6	<b>75.5</b>	96.8
w/o	GCE <sup>+</sup>	86.4	96.3	75.5	96.7
w/	GCE <sup>+</sup>	<b>89.3</b>	<b>97.1</b>	<b>76.0</b>	<b>96.7</b>

Table 2: Comparison of data augmentation results between VHDA with and without explicit goal tracking.

cant for less expressive state trackers (RNN) than the more expressive ones (GLAD<sup>+</sup> and GCE<sup>+</sup>). Even so, we observed varying degrees of improvements (zero to two percent in joint goal accuracy) even for the more expressive trackers, suggesting that GDA is effective regardless of downstream model expressiveness.

We observe larger improvement margins for inform-type dialog acts (or subsequently goals) from comparing performances between the dialog act types. This observation is because request-type dialog acts are generally more dependent on the user utterance in the same turn rather than requiring resolution of long-term dependencies, as illustrated in the dialog sample (Table 6). The observation supports our hypothesis that more diverse synthetic dialogs can benefit data augmentation by exploring unseen dialog dynamics.

Note that the goal tracking performances have relatively high variances due to the accumulative effect of tracking dialogs. However, as an additional benefit of employing GDA, we observe that synthetic dialogs help stabilize downstream tracking performances on DSTC2 and MultiWoZ-R datasets.

**Effects of Joint Goal Tracking.** Since user goals

DROP.	OBJ.	$\mathbf{z}^{(c)}$ -KL	WoZ2.0	
			GOAL	REQ
0.00	STD.	5.63	84.1±0.9	95.9±0.6
0.00	MIM	5.79	86.0±0.2	96.1±0.2
0.25	STD.	10.44	88.5±1.4	96.9±0.1
0.25	MIM	11.31	88.9±0.4	97.0±0.2
0.50	STD.	14.68	88.6±1.0	96.9±0.2
0.50	MIM	16.33	89.2±0.8	96.9±0.2
HIER.	STD.	14.34	88.2±1.0	97.1±0.2
HIER.	MIM	16.27	<b>89.3</b> ±0.4	<b>97.1</b> ±0.2

Table 3: Ablation studies on the training techniques using GCE<sup>+</sup> as the tracker. The effect of different dropout schemes and training objectives is quantified. MIM refers to mutual information maximization (§ 3.5).

can be inferred from turn-level inform-type dialog acts, it may seem redundant to incorporate goal modeling into our model. To verify its effectiveness, we train a variant of VHDA, where the model does not explicitly track goals. The results (Table 2) show that VDHA without explicit goal tracking suffers in joint goal accuracy but performs better in turn request accuracy for some instances. We conjecture that explicit goal tracking helps the model reinforce long-term dialog goals; however, the model does so in the minor expense of short-term state tracking (as evident from lower state tracking accuracy).

**Effects of Employing Training Techniques.** To demonstrate the effectiveness of the two proposed training techniques, we compare (1) the data augmentation results and (2) the KL-divergence between the posterior and prior of the dialog latents  $\mathbf{z}^{(c)}$  (Table 3). The results support our hypothesis that the proposed measures reduce the risk of inference collapse. We also confirm that exponentially-scaled dropouts are more or comparably effective at preventing posterior collapse than uniform

MODEL	WoZ2.0		DSTC2	
	ROUGE	ENT	ROUGE	ENT
VHCR <sup>a</sup>	0.476	0.193	0.680	0.153
VHDA <sup>b</sup> w/O GOAL	0.473	<b>0.195</b>	0.743	<b>0.162</b>
VHDA <sup>b</sup>	<b>0.499</b>	0.193	<b>0.781</b>	0.154

<sup>a</sup> (Park et al., 2018) <sup>b</sup> Ours

Table 4: Results on language quality and diversity evaluation.

MODEL	WoZ2.0		DSTC2	
	ACC	ENT	ACC	ENT
VHUS <sup>a</sup>	0.322	0.056	0.367	0.024
VHDA <sup>b</sup> w/O GT	0.408	0.079	0.460	0.034
VHDA <sup>b</sup>	<b>0.460</b>	<b>0.080</b>	<b>0.554</b>	<b>0.043</b>

<sup>a</sup> (Gür et al., 2018) <sup>b</sup> Ours

Table 5: Comparison of user simulation performances.

dropouts while generating more coherent samples (evident from higher data augmentation results).

### 4.3 Language Evaluation

To understand the effect of joint learning of various dialog features on language generation, we compare our model with a model that only learns linguistic features. Following the evaluation protocol from prior work (Wen et al., 2017; Bak and Oh, 2019), we use ROUGE-L F1-score (Lin, 2004) to evaluate the linguistic quality and utterance-level unigram cross-entropy (Serban et al., 2017) (regarding the training corpus distribution) to evaluate diversity. Table 4 shows that our model generates better and more diverse utterances than the previous strong baseline on conversation modeling. These results supports the idea that joint learning of dialog annotations improves utterance generation, thereby increasing the chance of generating novel samples that improve the downstream trackers.

### 4.4 User Simulation Evaluation

Simulating human participants has become a crucial feature for training dialog policy models using reinforcement learning and automatic evaluation of dialog systems (Asri et al., 2016). Although our model does not specialize in user simulation, our experiments show that the model outperforms the previous model (VHUS<sup>2</sup>) (Gür et al., 2018) in terms of accuracy and creativeness (diversity). We evaluate the user simulation quality using the pre-

<sup>2</sup>The previous model employs variational inference for contextualized sequence-to-sequence dialog act prediction.

SPKR.	UTTERANCE	GOAL	TURN ACT
1 User	i want to find a cheap restaurant in the north part of town .	inform(area=north) inform(price range=cheap)	inform(area=north) inform(price range=cheap)
2 Wizard	what food type are you looking for ?		request(slot=food)
3 User	any type of restaurant will be fine .	inform(area=north) inform(food=dontcare) inform(price range=cheap)	inform(food=dontcare)
4 Wizard	the <place> is a cheap indian restaurant in the north . would you like more information ?		
5 User	what is the number ?	inform(area=north) inform(food=dontcare) inform(price range=cheap)	request(slot=phone)
6 Wizard	<place> 's phone number is <number> . is there anything else i can help you with ?		
7 User	no thank you . goodbye .	inform(area=north) inform(food=dontcare) inform(price range=cheap)	

Table 6: A sample generated from the midpoint between two latent variables in the  $z^{(c)}$  space encoded from two anchor data points.

diction accuracy on the test sets and the diversity using the entropy<sup>3</sup> of predicted dialog act specifications (act-slot-value triples). We present the results in Table 5.

### 4.5 $z^{(c)}$ -interpolation

We conduct  $z^{(c)}$ -interpolation experiments to demonstrate that our model can generalize the dataset space and learn to decode plausible samples from unseen latent space. The generated sample (Table 6) shows that our model can maintain coherence while generalizing key dialog features, such as the user goal and the dialog length. As a specific example, given both anchors' user goals (food=mediterranean and food=indian, respectively)<sup>4</sup>, the generated midpoint between the two data points is a novel dialog with no specific food type (food=dontcare).

## 5 Conclusion

We proposed a novel hierarchical and recurrent VAE-based architecture to capture accurately the semantics of fully annotated goal-oriented dialog

<sup>3</sup>The entropy is calculated with respect to the training set distribution

<sup>4</sup>The supplementary material includes the full examples.



corpora. To reduce the risk of inference collapse while maximizing the generation quality, we directly modified the training objective and devised a technique to scale dropouts along the hierarchy. We showed that our proposed model VHDA was able to achieve significant improvements for various competitive dialog state trackers in diverse corpora through extensive experiments. With recent trends in goal-oriented dialog systems gravitating towards end-to-end approaches (Lei et al., 2018), we wish to explore a self-supervised model, which discriminatively generates samples that directly benefit the downstream models for the target task. We would also like to explore different implementations in line with recent advances in dialog models, especially using large-scale pre-trained language models.

## Acknowledgement

We thank Hyunsoo Cho for his help with implementations and Jihun Choi for the thoughtful feedback. We also gratefully acknowledge support from Adobe Inc. in the form of a generous gift to Seoul National University.

## References

- Antreas Antoniou, Amos Storkey, and Harrison Edwards. 2017. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*.
- Layla El Asri, Jing He, and Kaheer Suleman. 2016. A sequence-to-sequence model for user simulation in spoken dialogue systems. In *INTERSPEECH*, pages 1151–1155.
- JinYeong Bak and Alice Oh. 2019. Variational hierarchical user-based conversation model. In *EMNLP-IJCNLP*, pages 1941–1950.
- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. 2018. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *TACL*, 5:135–146.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *SIGLL*, pages 10–21.
- Paweł Budzianowski, TsungHsien Wen, BoHsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *EMNLP*, pages 5016–5026.
- Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. 2018. Isolating sources of disentanglement in variational autoencoders. In *NIPS*, pages 2610–2620.
- Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*.
- Chris Cremer, Xuechen Li, and David Duvenaud. 2018. Inference suboptimality in variational autoencoders. *arXiv preprint arXiv:1801.03558*.
- Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2018. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*.
- Laila Dybkjær and Wolfgang Minker. 2008. *Recent Trends in Discourse and Dialogue*, volume 39. Springer Science & Business Media.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2017. Data augmentation for low-resource neural machine translation. In *ACL*, pages 567–573.
- Xiaodong Gu, Kyunghyun Cho, JungWoo Ha, and Sunghun Kim. 2018. Dialogwae: Multimodal response generation with conditional wasserstein autoencoder. In *ICLR*.
- Izzeddin Gür, Dilek Hakkani-Tür, Gokhan Tür, and Pararth Shah. 2018. User modeling for task oriented dialogues. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 900–906. IEEE.
- Kazuma Hashimoto, Yoshimasa Tsuruoka, Richard Socher, et al. 2017. A joint many-task model: Growing a neural network for multiple nlp tasks. In *EMNLP*, pages 1923–1933.
- Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534*.
- Matthew Henderson, Blaise Thomson, and Jason D Williams. 2014. The second dialog state tracking challenge. In *SIGDIAL*, pages 263–272.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.

- Matthew D Hoffman and Matthew J Johnson. 2016. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *the NIPS Workshop in Advances in Approximate Bayesian Inference*, volume 1.
- Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. 2020. A simple language model for task-oriented dialogue. *arXiv preprint arXiv:2005.00796*.
- Yutai Hou, Yijia Liu, Wanxiang Che, and Ting Liu. 2018. Sequence-to-sequence data augmentation for dialogue language understanding. In *ICCL*, pages 1234–1245.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *ICML*, pages 1587–1596. JMLR. org.
- Sungdong Kim, Sohee Yang, Gyuwan Kim, and Sang-Woo Lee. 2019. Efficient dialogue state tracking by selectively overwriting memory. *arXiv preprint arXiv:1911.03906*.
- Yoon Kim, Sam Wiseman, Andrew Miller, David Sontag, and Alexander Rush. 2018. Semi-amortized variational autoencoders. In *International Conference on Machine Learning*, pages 2678–2687.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. 2015. Audio augmentation for speech recognition. In *ISCA*.
- Gakuto Kurata, Bing Xiang, Bowen Zhou, and Mo Yu. 2016. Leveraging sentence-level information with encoder lstm for semantic slot filling. In *EMNLP*, pages 2077–2083.
- Wenqiang Lei, Xisen Jin, MinYen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *ACL*, pages 1437–1447.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- Ramesh Manuvinakurike, Jacqueline Brixey, Trung Bui, Walter Chang, Ron Artstein, and Kallirroi Georgila. 2018. Dialedit: Annotations for spoken conversational image editing. In *Proceedings 14th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, pages 1–9.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, TsungHsien Wen, Blaise Thomson, and Steve Young. 2017. Neural belief tracker: Data-driven dialogue state tracking. In *ACL*, pages 1777–1788.
- Elnaz Nouri and Ehsan Hosseini-Asl. 2018. Toward scalable neural dialogue state tracking model. *arXiv preprint arXiv:1812.00899*.
- Yookoon Park, Jaemin Cho, and Gunhee Kim. 2018. A hierarchical latent structure for variational conversation modeling. In *NAACL:HLT*, pages 1792–1801.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Ali Razavi, Aäron van den Oord, Ben Poole, and Oriol Vinyals. 2018. Preventing posterior collapse with delta-vaes. In *ICLR*.
- Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-based user simulation for bootstrapping a pomdp dialogue system. In *HLT: NAACL*, pages 149–152.
- John R Searle, Ferenc Kiefer, Manfred Bierwisch, et al. 1980. *Speech act theory and pragmatics*, volume 10. Springer.
- Iulian Vlad Serban, Alessandro Sordani, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *AAAI*, pages 3295–3301. AAAI Press.
- Youhyun Shin, Kang Min Yoo, and Sanggoo Lee. 2019. Utterance generation with variational auto-encoder for slot filling in spoken language understanding. *IEEE Signal Processing Letters*, 26(3):505–509.
- Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):60.
- Blaise Thomson and Steve Young. 2010. Bayesian update of dialogue state: A pomdp framework for spoken dialogue systems. *Computer Speech & Language*, 24(4):562–588.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. 2018. **Wasserstein auto-encoders**. In *International Conference on Learning Representations*.
- Toan Tran, Trung Pham, Gustavo Carneiro, Lyle Palmer, and Ian Reid. 2017. A bayesian data augmentation approach for learning deep models. In *NeurIPS*, pages 2797–2806.
- Zhuoran Wang and Oliver Lemon. 2013. A simple and generic belief tracking mechanism for the dialog state tracking challenge: On the believability of observed information. In *SIGDIAL*, pages 423–432.

- Jason W Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*, pages 438–449.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. *arXiv preprint arXiv:1905.08743*.
- Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, and Qun Liu. 2019. Dialog state tracking with reinforced data augmentation. *arXiv preprint arXiv:1908.07795*.
- Kang Min Yoo, Youhyun Shin, and Sangwoo Lee. 2019. Data augmentation for spoken language understanding via joint variational generation. In *AAAI*, volume 33, pages 7402–7409.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NeurIPS*, pages 649–657.
- Shengjia Zhao, Jiaming Song, and Stefano Ermon. 2017. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2018. Global-locally self-attentive encoder for dialogue state tracking. In *ACL*, pages 1458–1467.
- Lukas Zilka and Filip Jurcicek. 2015. Incremental lstm-based dialog state tracker. In *2015 IEEE Workshop on ASRU*, pages 757–762. IEEE.

## Appendix A Mutual Information Maximization for Mitigating Inference Collapse

During the training of VAEs, inference collapse occurs when the model converges to a local optimum where the approximate posterior  $q_\phi(\mathbf{z} | \mathbf{x})$  collapses to the prior  $p(\mathbf{z})$ , indicating the vanishment of the encoder network due to the decoder’s negligence of the encoder signals. Quantifying, diagnosing, and devising a mitigation technique for the inference collapse phenomenon have been studied extensively in the past (Chen et al., 2016; Zhao et al., 2017; Cremer et al., 2018; Razavi et al., 2018; He et al., 2019). However, current approaches for mitigating inference collapse are limited to significant modifications to the existing VAE framework (He et al., 2019; Kim et al., 2018) or limited to specific architectural designs (Razavi et al., 2018). Current approaches do not work well on our model due to the complexity of our VAE structure. Instead, we employ a relatively simple technique that directly modifies the VAE objective. By doing so, we mitigate any significant changes to the main VAE framework while achieving satisfactory results on inference collapse mitigation. Though not covered in this paper, our method has applications in other VAE structures. In this appendix, we wish to delve more in-depth into the intuitions and detailed implementation of our approach.

**Motivation.** As first noted by Hoffman and Johnson (2016) (and subsequently utilized by (Zhao et al., 2017; Chen et al., 2018)), the KL-divergence term of the ELBO objective can be decomposed into two terms: (1) the KL-divergence between the aggregate posterior and the prior and (2) the mutual information between the latent variables and the data:

$$\begin{aligned} \mathbb{E}_{p_d}[D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z}))] = \\ D_{\text{KL}}(q_\phi(\mathbf{z}) || p(\mathbf{z})) + I_{q_\phi}(\mathbf{x}; \mathbf{z}) \end{aligned} \quad (3)$$

where  $p_d$  is the empirical distribution of data and the aggregate posterior  $q_\phi(\mathbf{z})$  is obtained by marginalizing the approximate posterior using the empirical distribution:

$$q_\phi(\mathbf{z}) = \mathbb{E}_{\mathbf{x} \sim p_d}[q_\phi(\mathbf{z} | \mathbf{x})]. \quad (4)$$

Using the definition of inference collapse, we can deduce that the KL-divergence term  $D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z}))$  is zero during inference collapse. This

fact implies that both decomposed terms in Equation 3 must be zero since both terms are non-negative.

Our preliminary studies show an interesting pattern in the KL-divergence term and its decomposed terms during basic training (training without inference-collapse treatments) (Figure 2). We observe that the KL-divergence of the aggregate posterior term vanishes earlier than the mutual information does. We also observe that the mutual information term, which represents the encoder effectiveness, vanishes eventually. This collapse happens after the KL-divergence cannot be minimized without sacrificing the encoder’s expressiveness. Note that optimization of the ELBO objective minimizes the ELBO’s KL-divergence term and its underlying terms, one of which is directly related to the encoder health. Although the reconstruction term in the ELBO encourages maximization of the mutual information, the autoregressive property of the decoder and the complexity of the reconstruction loss “dilutes” the goal of maximizing mutual information. Hence, to minimize inference collapse, we propose a modified VAE objective that explicitly maximizes the mutual information between the latents and the data by “canceling” out the mutual information term in the KL-divergence<sup>5</sup>:

$$\begin{aligned} \mathcal{L}_{\text{VHDA}} = & \mathbb{E}_{p_d}[\mathbb{E}_{q_\phi}[\log p_\theta(\mathbf{c} | \mathbf{z})]] \\ & - \mathbb{E}_{p_d}[D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{c}) || p(\mathbf{z}))] \\ & + I_{q_\phi}(\mathbf{c}; \mathbf{z}). \end{aligned} \quad (5)$$

Note that some notations (expectation over the empirical distribution) have been omitted in the main paper for clarity.

**Relation to Prior Work.** Our approach is related to previous work on manipulating the VAE objective for customizing the VAE behavior (Zhao et al., 2017; Chen et al., 2018). It can also be thought of as a special case of Wasserstein Autoencoders (Tolstikhin et al., 2018) Although not all related works were original proposed to directly combat inference collapse, our approach can be considered a special case of InfoVAE (Zhao et al., 2017) and  $\beta$ -TCVAE (Chen et al., 2018). Specifically, Zhao et al. (2017) proposed a modified VAE objective as

<sup>5</sup>On a side note, we did not observe any “lag” in the inference network, as described by He et al. (2019). This observation is evident from the sustained mutual information level throughout the training session (Figure 2). Hence we did not employ the recently proposed method.






Figure 2: Failed training behavior.

follows:

$$\begin{aligned} \mathcal{L}_{\text{InfoVAE}} = & \mathbb{E}_{p_d}[\mathbb{E}_{q_\phi}[\log p_\theta(\mathbf{x} | \mathbf{z})]] \\ & - (1 - \alpha)\mathbb{E}_{p_d}[D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z}))] \\ & - (\alpha + \lambda - 1)D_{\text{KL}}(q_\phi(\mathbf{z}) || p(\mathbf{z})). \end{aligned} \quad (6)$$

Rearranging the equation, we can express the same objective related to the mutual information:

$$\begin{aligned} \mathcal{L}_{\text{InfoVAE}} = & \mathbb{E}_{p_d}[\mathbb{E}_{q_\phi}[\log p_\theta(\mathbf{x} | \mathbf{z})]] \\ & - \lambda D_{\text{KL}}(q_\phi(\mathbf{z}) || p(\mathbf{z})) \\ & - (1 - \alpha)I_{q_\phi}(\mathbf{x}; \mathbf{z}). \end{aligned} \quad (7)$$

Hence, our method is a special case of InfoVAE where  $\alpha = 1$  and  $\lambda = 1$ . Meanwhile, [Chen et al. \(2018\)](#) proposed an extended modification to  $\beta$ -VAE ([Higgins et al., 2017](#)) to further decompose the KL-divergence of the aggregate posterior in terms of latent correlation:

$$\begin{aligned} \mathcal{L}_{\text{InfoVAE}} = & \mathbb{E}_{p_d}[\mathbb{E}_{q_\phi}[\log p_\theta(\mathbf{x} | \mathbf{z})]] \\ & - \alpha I_{q_\phi}(\mathbf{x}; \mathbf{z}) \\ & - \beta D_{\text{KL}}(q_\phi(\mathbf{z}) || \sum_i q_\phi(z_i)) \\ & - \gamma \sum_i D_{\text{KL}}(q_\phi(z_i) || p(z_i)). \end{aligned} \quad (8)$$

In the equation above, our approach corresponds the case where  $\alpha = 0$  and  $\beta = \gamma = 1$ .

**Mutual Information Estimation.** We can estimate the mutual information between the latents and the data under the empirical distribution of  $\mathbf{x}$

using Monte Carlo sampling. However, this estimation method is known to be biased ([Belghazi et al., 2018](#)). Despite recent advances in MI estimation techniques, we find that our unparameterized method is sufficient for achieving inference collapse mitigation and probing.:

The equation for estimating the mutual information is shown in [Equation 9](#). where  $\mathbf{x}$  is sampled from the empirical distribution of the dataset and  $N$ ,  $M$  and  $L$  are hyperparameters. In practice, the estimation is performed over the data samples in a mini-batch for computational efficiency. Given a mini-batch of size  $N$ , we further approximate the estimation by sampling the latent variables  $\mathbf{z}$  once for each data point ( $M = 1$ ) ([Equation 10](#)).

We visualize the variance in our mutual information estimation method in [Figure 3](#).

## Appendix B Architectural Diagram

We include a more detailed architectural diagram ([Figure 4](#)) depicting the latent variables and the model inference, which we could not illustrate in [Figure 1](#) due to space constraints. Note that the orange crosses denote decoder dropouts. The figure also illustrates the hierarchically-scaled dropout scheme, motivated by the need to minimize information loss while discouraging the decoders from relying on training signals, leading to exposure bias.

$$I_{q_\phi}(\mathbf{x}, \mathbf{z}) = \mathbb{E}_{p_d} [D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{x}) \| q_\phi(\mathbf{z}))]$$

$$\approx \frac{1}{NM} \sum_i^N \sum_j^M \left( \log q_\phi(\mathbf{z}_j | \mathbf{x}_i) - \log \sum_k^L q_\phi(\mathbf{z}_j | \mathbf{x}_k) + \log L \right) \quad (9)$$

$$I_{q_\phi}(\mathbf{x}, \mathbf{z}) \approx \frac{1}{N} \sum_i^N \left[ \log q_\phi(\mathbf{z} | \mathbf{x}_i) - \log \sum_j^N q_\phi(\mathbf{z} | \mathbf{x}_j) + \log N \right]_{\mathbf{z} \sim q_\phi(\mathbf{z} | \mathbf{x}_i)} \quad (10)$$

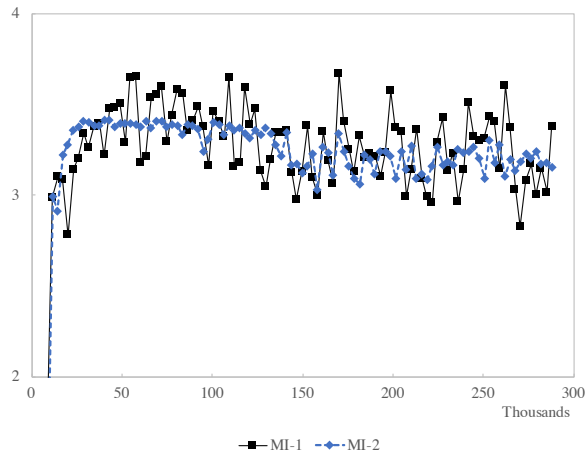


Figure 3: Estimation of the mutual information over the course of training. MI-2 corresponds to our approach. MI-1 is derived from the Monte Carlo estimation of  $D_{\text{KL}}(q_\phi(\mathbf{z}) \| p(\mathbf{z}))$  (not described). Our approach results in less variance in the MI estimation.

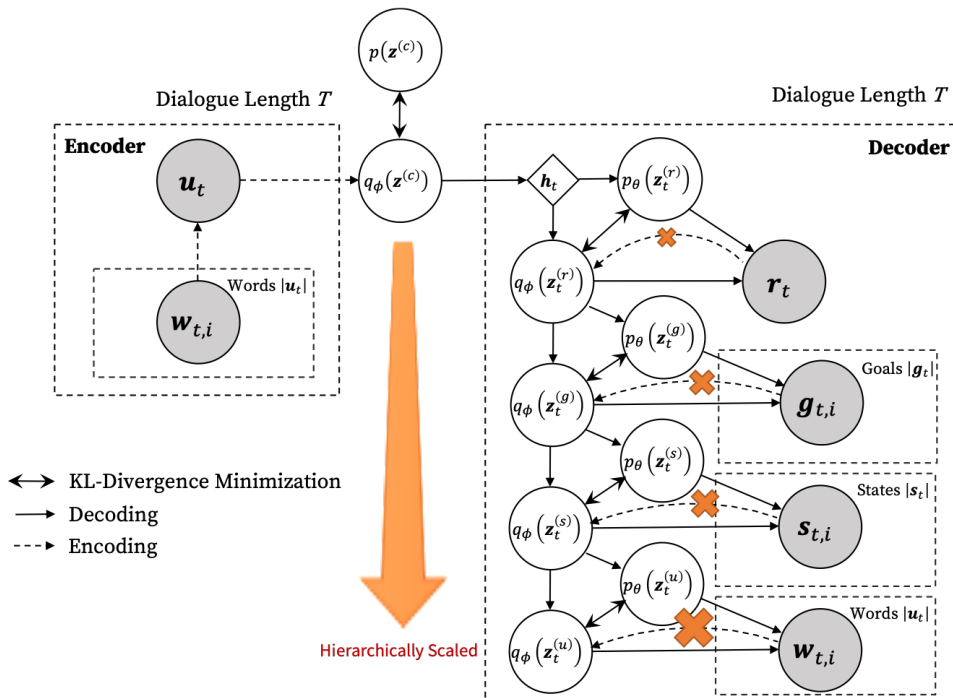


Figure 4: The architectural diagram.

## Appendix C Full Results of Data Augmentation

The full data augmentation results are shown below, including the statistics. Note that generative data augmentation also has the effect of reducing the variance of the downstream models.

GDA	MODEL	WoZ2.0			DSTC2			MWoZ-R			MWoZ-H			DIALEdit		
		GOAL	REQ		GOAL	REQ		GOAL	INF		GOAL	INF		GOAL	REQ	
-	RNN	74.5±0.8	96.1±0.3	69.7±7.2	96.0±0.4	43.7±8.7	69.4±5.7	25.7±4.1	55.6±2.3	35.8±3.1	96.6±0.5					
VHDA	RNN	<b>78.7±2.1<sup>†</sup></b>	<b>96.7±0.1<sup>‡</sup></b>	<b>74.2±0.9<sup>†</sup></b>	<b>97.0±0.2<sup>‡</sup></b>	<b>49.6±3.1<sup>†</sup></b>	<b>73.4±1.8<sup>†</sup></b>	<b>31.0±5.0<sup>†</sup></b>	<b>59.7±3.1<sup>†</sup></b>	<b>36.4±1.4<sup>†</sup></b>	<b>96.8±0.1</b>					
-	GLAD <sup>+</sup>	87.8±0.8	<b>96.8±0.3</b>	74.5±0.5	96.4±0.2	58.9±2.5	76.3±1.4	33.4±2.4	58.9±1.5	35.9±1.0	96.7±0.3					
VHDA	GLAD <sup>+</sup>	<b>88.4±0.3</b>	96.6±0.2	<b>75.5±0.5<sup>‡</sup></b>	<b>96.8±0.5<sup>†</sup></b>	<b>61.5±2.4<sup>†</sup></b>	<b>77.4±2.0</b>	<b>37.8±2.2<sup>‡</sup></b>	<b>61.3±1.0<sup>†</sup></b>	<b>37.1±1.1<sup>†</sup></b>	<b>96.8±0.4</b>					
-	GCE <sup>+</sup>	88.3±0.7	97.0±0.2	74.8±0.6	96.3±0.2	60.5±3.4	76.7±1.2	36.5±2.4	61.0±1.2	36.1±1.3	96.6±0.4					
VHDA	GCE <sup>+</sup>	<b>89.3±0.4<sup>‡</sup></b>	<b>97.1±0.2</b>	<b>76.0±0.2<sup>‡</sup></b>	<b>96.7±0.4<sup>†</sup></b>	<b>63.3±3.9</b>	<b>77.2±3.3</b>	<b>38.3±4.1</b>	<b>63.1±1.4<sup>†</sup></b>	<b>37.6±2.1<sup>†</sup></b>	<b>96.8±0.4</b>					

<sup>†</sup>  $p < 0.1$

<sup>‡</sup>  $p < 0.01$

Table 7: The full results of data augmentation, including the standard deviations of 9 repeated runs.

## Appendix D Exhibits of Synthetic Samples

This section describes the method we use to sample synthetic data points from our model’s posterior and presents some synthetic samples generated from our model using the described technique.

We use *ancestral sampling* (He et al., 2019), or the *posterior sampling* technique (Yoo et al., 2019), to sample data points from the empirical distribution of the latent space. Specifically, we choose an anchor data point from the dialog dataset:  $\mathbf{c} \sim p_d(\mathbf{c})$ , where  $p_d$  is the empirical distribution of goal-oriented dialogs. Then, we sample a set of latent variables  $\mathbf{z}^{(c)}$  from the encoded distribution of  $\mathbf{c}$ :  $\mathbf{z}^{(c)} \sim q_\phi(\mathbf{z}^{(c)} | \mathbf{c})$ . Next, we decode a sample  $\mathbf{c}'$  that maximizes the log-likelihood for each sampled conversational latents:

$$\mathbf{c}' = \arg \max_{\mathbf{c}} p_\theta(\mathbf{c} | \mathbf{z}^{(c)}).$$

We use these samples to augment the original dataset. Also, we fix the ratio of the synthetic dataset to the original dataset to 1. In our experiments, we observe that all of the synthetic samples generated via ancestral sampling are mostly coherent and, most importantly, novel, i.e., each synthetic data point is somehow different from the original anchor point (e.g., variations in utterances, dialog-level semantics, or sometimes annotation errors).

In the following tables, we showcase few dialog samples from our augmentation datasets. The tables present the generated samples along with their reference dialog samples.

SPEAKER	UTTERANCE	GOAL	TURN ACT
ANCHOR (REAL)			
1	User	i am looking for a panasian restaurant in the south side of town . if there are n't any maybe chinese . i need an address and price	inform(area=south) inform(food=panasian) request(slot=price range) request(slot=address)
2	Wizard	there is an expensive and a cheap chinese restaurant in the south . which would you prefer ?	request(slot=price range)
3	User	let 's try cheap chinese restaurant . can i get an address ?	inform(area=south) inform(food=chinese) inform(price range=cheap) request(slot=address)
4	Wizard	of course it 's <location>	
5	User	thank you goodbye .	inform(area=south) inform(food=chinese) inform(price range=cheap)
POINTWISE POSTERIOR SAMPLE (GENERATED)			
1	User	i 'm looking for a panasian restaurant in the south side of town . if there are n't any maybe chinese . i need an address and price	inform(food=panasian) inform(area=south) request(slot=price range)
2	Wizard	there are no cheap restaurants serving restaurants i have a seafood the the number is some other available	
3	User	how about thai	inform(food=thai) inform(price range=cheap) request(slot=address)
4	Wizard	we 's <place> on <location>	
5	User	thank you very much .	inform(food=thai) inform(price range=cheap)



SPEAKER	UTTERANCE	GOAL	TURN ACT
ANCHOR (REAL)			
1	User	i need the address of a gastropub in town .	inform(food=gastropub) request(slot=address)
2	Wizard	which part of town ?	request(slot=area)
3	User	does n't matter .	inform(food=gastropub) inform(area=dont care)
4	Wizard	would you prefer moderate or expensive pricing ?	request(slot=price range)
5	User	moderate please .	inform(food=gastropub) inform(area=dont care) inform(price range=moderate)
6	Wizard	i have found one results that matches your criteria the restaurant the <place> is a gastropub located at <location> some code as the price range is moderate	
7	User	are there any others in that price range ?	inform(food=gastropub) inform(area=dont care) inform(price range=moderate)
8	Wizard	unfortunately there are not sorry	
9	User	hello i am looking for a restaurant that serves gastropub food in any area can you help me ?	inform(food=gastropub) inform(area=dont care) inform(price range=moderate)
10	Wizard	sure would you prefer expensive or moderately priced ?	request(slot=price range)
11	User	thank you goodbye	inform(food=gastropub) inform(area=dont care) inform(price range=moderate)
POINTWISE POSTERIOR SAMPLE (GENERATED)			
1	User	i need the address of a gastropub in town .	inform(food=gastropub) inform(food=gastropub)
2	Wizard	i have many options . would you prefer centre or east ?	request(slot=area)
3	User	does n't matter .	inform(food=gastropub) inform(area=dont care) inform(area=center)
4	Wizard	there are three gastropub restaurants listed . one is in the east part of town and the rest are in the centre .	request(slot=price range)
5	User	i do n't care	inform(food=gastropub) inform(price range=moderate) inform(area=dont care)
6	Wizard	i found <place> . results that matches your criteria the restaurant the <place> is a gastropub located at <location> some <code> as the price range is moderate	
7	User	are there any others in that price range ?	inform(food=gastropub) inform(price range=moderate) inform(area=dont care)
8	Wizard	in that actually not sorry	
9	User	hello i am looking for a restaurant that serves gastropub food in any area can you help me ?	inform(food=gastropub) inform(price range=moderate) inform(area=dont care)

SPEAKER	UTTERANCE	GOAL	TURN ACT
ANCHOR (REAL)			
1	User	i 'm looking for a cheap restaurant in the west part of town .	inform(area=west) inform(price range=cheap)
2	Wizard	i found a vietnamese and italian cheap restaurant in the west side of town . would you like the phone number or address of either ?	request(slot=phone) request(slot=address)
3	User	yes please .	inform(area=west) inform(price range=cheap)
4	Wizard	<place> is the italian restaurant located at <location> . its phone number is <numeric> . <place> is the vietnamese restaurant located at <location> . its phone number is	
5	User	thank you .	inform(area=west) inform(price range=cheap)
6	Wizard	you 're welcome	
7	User	goodbye .	inform(area=west) inform(price range=cheap)
POINTWISE POSTERIOR SAMPLE 1 (GENERATED)			
1	User	i 'm looking for a cheap restaurant in the west part of town .	inform(price range=cheap) inform(area=west)
2	Wizard	there is a cheap restaurant in the west part of town . would you like their address and location ?	request(slot=phone) request(slot=address)
3	User	yes please .	inform(area=west) inform(area=north) inform(price range=cheap)
4	Wizard	<place> is the italian restaurant .	
5	User	thank you very much goodbye .	inform(area=north) inform(price range=cheap)
POINTWISE POSTERIOR SAMPLE 2 (GENERATED)			
1	User	i want a cheap restaurant on the west side .	inform(price range=cheap) inform(area=west)
2	Wizard	<place> is a restaurant that matches your choice in the west .	
3	User	<place> the phone and the address ?	inform(food=vietnamese) inform(price range=cheap) inform(area=west)
4	Wizard	<place> 's phone number is <numeric>	
5	User	thank you that will do .	inform(food=vietnamese) inform(price range=cheap) inform(area=west)

## Appendix E $z^{(c)}$ Interpolation Results (Including Both Anchors)

Visualizing samples from a linear interpolation of two points in the latent space (Bowman et al., 2016) is a popular way to showcase the generative capability of VAEs. Given two dialog samples  $c_1$  and  $c_2$ , we map the data points onto the conversational latent space to obtain  $z_1^{(c)}$  and  $z_2^{(c)}$ . Multiple equidistant samples  $z'_1, \dots, z'_N$  are selected from the linear interpolation between the two points:  $z'_n = z_1^{(c)} + n(z_2^{(c)} - z_1^{(c)})/N$ . Likelihood-maximizing samples  $x'_1, \dots, x'_N$  are chosen from the model posteriors given the intermediate latent samples.

SPEAKER	UTTERANCE	GOAL	TURN ACT
ANCHOR 1 (REAL)			
1	User	i 'm looking for a mediterranean place for any price . what is the phone and postcode ?	inform(food=mediterranean) inform(price=dont care) request(slot=phone) request(slot=postcode)
2	Wizard	i found a few places . the first is <place> with a phone number of <number> and postcode of <postcode>	
3	User	That will be fine . thank you .	inform(food=mediterranean) inform(price=dont care)
MIDPOINT 50% (GENERATED)			
1	User	i want to find a cheap restaurant in the north part of town .	inform(area=north) inform(price range=cheap)
2	Wizard	what food type are you looking for ?	request(slot=food)
3	User	any type of restaurant will be fine .	inform(area=north) inform(food=dontcare) inform(price range=cheap)
4	Wizard	the <place> is a cheap indian restaurant in the north . would you like more information ?	
5	User	what is the number ?	inform(area=north) inform(food=dontcare) inform(price range=cheap)
6	Wizard	<place> 's phone number is <number> . is there anything else i can help you with ?	
7	User	no thank you . goodbye .	inform(area=north) inform(food=dontcare) inform(price range=cheap)
ANCHOR 2 (REAL)			
1	User	i am looking for a cheap restaurant in the north part of town .	inform(area=north) inform(price range=cheap)
2	Wizard	there are two restaurants that fit your criteria would you prefer italian or indian food ?	request(slot=food)
3	User	let s try indian please	inform(area=north) inform(price range=cheap) inform(food=indian)
4	Wizard	<name> serves indian food in the cheap price range and in the north part of town . is there anything else i can help you with ?	
5	User	what is the name of the italian restaurant ?	inform(area=north) inform(price range=cheap) inform(food=indian)
6	Wizard	<name>	
7	User	what is the address and phone number ?	inform(area=north) inform(price range=cheap) inform(food=indian)
8	Wizard	the address for <name> is <address> and the phone number is <phone> .	
9	User	thanks so much .	inform(area=north) inform(price range=cheap) inform(food=indian)

SPEAKER	UTTERANCE	GOAL	TURN ACT
ANCHOR 1 (REAL)			
1	User	hi i 'm looking for a moderately priced restaurant in the south part of town .	inform(area=south) inform(price range=moderate)
2	Wizard	the <place> <location> is moderately priced and in the south part of town . would you like their location ?	request(slot=address)
3	User	yes . i would like the location and the phone number please .	inform(area=south) inform(price range=moderate)
4	Wizard	the address of <place> <location> is <location> and the phone number is <numeric> .	request(slot=phone) request(slot=address)
5	User	thank you goodbye .	inform(area=south) inform(price range=moderate)
30% (GENERATED)			
1	User	i am looking for some seafood what can you tell me ?	inform(area=dont care) inform(food=seafood)
2	Wizard	<place> restaurant bar serves mexican food in the south part of town . would you like their location ?	inform(area=dont care) request(slot=address)
3	User	yes i 'd like the address phone number and postcode please .	inform(food=lebanese) inform(food=seafood)
4	Wizard	<place> is located at <location> cost the phone number is <numeric> .	request(slot=address) request(slot=phone)
5	User	thank you goodbye .	inform(food=seafood) inform(area=dont care)
70% (GENERATED)			
1	User	i would like to find a restaurant in the east part of town that serves gastropub food .	inform(food=mexican) inform(food=mexican)
2	Wizard	<place> restaurant bar serves mexican food in the south part of town . would you like their location ?	request(slot=address)
3	User	yes i 'd like the address phone number and postcode please .	inform(food=mexican) request(slot=address) request(slot=postcode) request(slot=phone)
4	Wizard	<place> restaurant bar is located at <location> . the postal code is some code and the phone number is <numeric> .	
5	User	thank you goodbye .	inform(food=mexican)
ANCHOR 2 (REAL)			
1	User	i want to find a restaurant in any part of town and serves malaysian food .	inform(area=dont care) inform(food=malaysian)
2	Wizard	there are no malaysian restaurants . would you like something different ?	inform(area=dont care) inform(food=malaysian)
3	User	north american please . give me their price range and their address and phone number please .	inform(food=north american) request(slot=phone) request(slot=price range) request(slot=address)
4	Wizard	<place> is in the expensive price range their phone number is <numeric> and their address is <location>	
5	User	thank you goodbye	inform(area=dont care) inform(food=north american)