

Out-of-Domain Detection for Natural Language Understanding in Dialog Systems

Yinhe Zheng¹, Guanyi Chen², and Minlie Huang³, *Member, IEEE*

Abstract—Natural Language Understanding (NLU) is a vital component of dialogue systems, and its ability to detect Out-of-Domain (OOD) inputs is critical in practical applications, since the acceptance of the OOD input that is unsupported by the current system may lead to catastrophic failure. However, most existing OOD detection methods rely heavily on manually labeled OOD samples and cannot take full advantage of unlabeled data. This limits the feasibility of these models in practical applications. In this paper, we propose a novel model to generate high-quality pseudo OOD samples that are akin to In-Domain (IND) input utterances and thereby improves the performance of OOD detection. To this end, an autoencoder is trained to map an input utterance into a latent code. Moreover, the codes of IND and OOD samples are trained to be indistinguishable by utilizing a generative adversarial network. To provide more supervision signals, an auxiliary classifier is introduced to regularize the generated OOD samples to have indistinguishable intent labels. Experiments show that these pseudo OOD samples generated by our model can be used to effectively improve OOD detection in NLU. Besides, we also demonstrate that the effectiveness of these pseudo OOD data can be further improved by efficiently utilizing unlabeled data.

Index Terms—Natural language understanding, out-of-domain detection, dialogue system, text classification.

I. INTRODUCTION

NATURAL Language Understanding (NLU) in dialog systems such as task-oriented dialog systems and intelligent personal assistants is vital for understanding users' input to make effective human-machine interaction. An NLU module maps unstructured text inputs to structured dialog acts and has a crucial influence on the downstream processing pipelines of a dialog system. Therefore, the reliability of NLU becomes a precursor to the success of dialog systems. Recently, various deep neural

network-based NLU models are proposed, and some of these models have been applied in real-world applications [1]–[3].

Most existing neural NLU modules are built by following a *closed-world* assumption [4], [5], i.e., the data used in training and testing phrase are drawn from the same distribution. However, such an assumption is commonly violated in practical systems that are deployed in a dynamic or open environment. Specifically, practical NLU modules often encounter out-of-domain (OOD) inputs that are not supported by the system and thus not observed in the training data. Wrongly accepting these inputs and executing undesired commands may trigger catastrophic failures, particularly in risk-sensitive applications where safety is the top priority, such as robots or self-driving cars. In order to address this issue, a more realistic assumption of *open-world* [4], [5] has been proposed. An NLU system built under this assumption should be able to not only correctly analyze in-domain (IND) inputs but also reliably reject OOD inputs that are not supported by the system.

Various methods have been proposed to improve the OOD detection performance of neural NLU models [6]–[10], and most of them follow a threshold-based protocol. Specifically, a detection score is computed for each input, and then a threshold is selected using a validation set. The inputs whose scores are lower than the threshold are considered to be OOD inputs and then rejected. A simple yet efficient approach is to use the maximum value of the Softmax output as the detection score [11], which has been demonstrated to work surprisingly well on the image classification tasks [12], and thus has been applied in many state-of-the-art systems [12], [13]. Further developments in this direction propose to add an extra entropy regularization (ER) term in the training objective, and significant performance improvement for OOD detection is reported when this ER term is optimized using a set of OOD data [14], [15].

However, collecting large-scale OOD data is usually difficult and expensive in practice, especially when dealing with the ever-changing *open-world* environment. This limits the feasibility of the ER technique in practical applications. To address this issue, some studies [14], [16] have proposed to generate pseudo OOD samples with a generative adversarial network (GAN) [17]. These generated samples can be used to optimize the ER term and thus improve the OOD detection performance. However, existing approaches only work in continuous space (such as generating OOD images or continuous feature vectors), whereas the input to NLU modules is usually a sequence of discrete tokens. It is yet to be explored to generate pseudo OOD samples in discrete spaces (like natural language) that can be

Manuscript received October 27, 2019; revised February 28, 2020; accepted March 17, 2020. Date of publication April 1, 2020; date of current version April 23, 2020. This work was supported by the National Science Foundation of China under Grant 61936010 and regular project with under Grant 61876096 and in part by the National Key R&D Program of China under Grant 2018YFC0830200. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Carlos Busso. (*Corresponding author: Minlie Huang.*)

Yinhe Zheng is with the Samsung Research China - Beijing (SRC-B), Beijing 100102, China, and also with the Tsinghua University, Beijing 100084, China (e-mail: yh.zheng@samsung.com; zhengyinhe1@163.com).

Guanyi Chen is with the Department of Information and Computing Sciences, Utrecht University, 3584, CC, Utrecht, The Netherlands (e-mail: g.chen@uu.nl).

Minlie Huang is with the Institute for Artificial Intelligence, State Key Lab of Intelligent Technology and Systems, Beijing National Research Center for Information Science and Technology, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (e-mail: aihuang@tsinghua.edu.cn).

Digital Object Identifier 10.1109/TASLP.2020.2983593

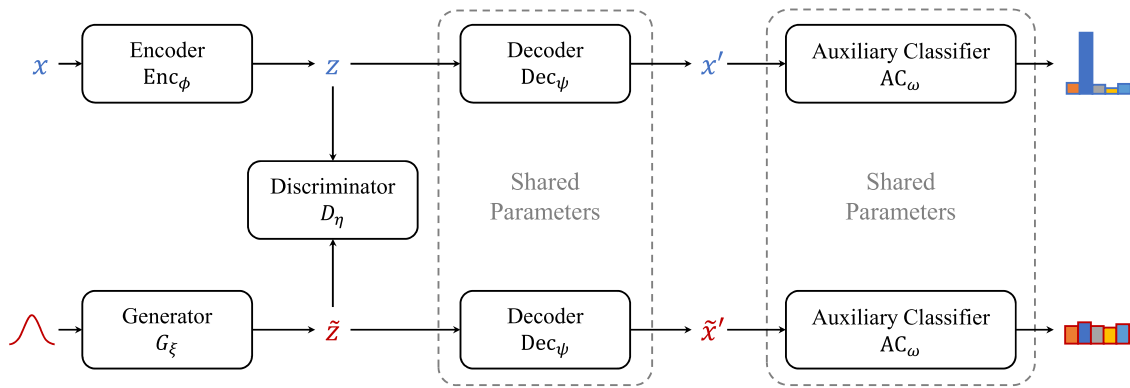


Fig. 1. Overall architecture of the pseudo OOD sample generation (POG) model. An encoder Enc_ϕ transforms an input utterance x to a latent code z . A decoder Dec_ψ reconstructs x' from z . A generator G_ξ is built to map a Gaussian noise to a fake latent code \tilde{z} and a discriminator D_η distinguishes the fake code \tilde{z} and the real code z with an adversarial training process. An auxiliary classifier AC_ω is trained to predict the correct label associated with x' , and G_ξ is regularized by the gradients derived from AC_ω which enforces utterance \tilde{x}' generated from \tilde{z} to follow an uniform distribution.

effectively utilized to improve the OOD detection performance. Moreover, unlabeled data (i.e., a mixture of IND and OOD samples) are usually easier to obtain in practical applications (e.g., through user logs), but rarely utilized in existing OOD detection methods. It is attractive to take advantage of these unlabeled data to improve the OOD detection performance since we can expect to gain some prior knowledge about the testing OOD distribution through these unlabeled data.

In this paper, we study how generated pseudo OOD samples and unlabeled data can facilitate OOD detection in NLU systems. We follow the simple and efficient approach to add an ER term in the training objective. However, we focus more on the way of generating high-quality pseudo OOD samples to effectively optimize this term. To this end, we propose a novel pseudo OOD sample generation model (POG) (depicted in Fig. 1).

The proposed model POG consists of three components: 1) a reconstruction module, 2) an adversarial generation module, and 3) an auxiliary classifier. In the reconstruction module, an encoder maps a text input into a latent code, and a decoder reconstructs the text from the latent code. A generator is then trained to produce fake latent codes, and a discriminator is trained to distinguish these fake codes from the real ones with an adversarial training process. In order to provide more supervision signals, an auxiliary classifier is further introduced to predict the correct labels associated with the reconstructed samples and to regularize the generated OOD samples to have indistinguishable labels. Experiments show that the OOD samples generated using our model can effectively improve the performance of OOD detection. We also demonstrate that unlabeled data can be used to train the reconstruction module and thus boost the effectiveness of the generated pseudo OOD samples.

To summarize, our contributions are in three folds:

- 1) We propose a novel model to generate pseudo OOD samples. The model consists of an autoencoder, an adversarial training component, and an auxiliary classifier. The generated samples can be used to effectively improve the OOD detection performance of NLU by optimizing the entropy regularization (ER) term in the training stage.

- 2) The proposed model can take advantage of unlabeled data to improve the effectiveness of the generated OOD samples. This makes our model more suitable for practical applications.
- 3) We evaluate the model on two datasets. Results show that our model can significantly outperform other competitive baselines for OOD detection.

II. RELATED WORK

The problem of OOD detection has been investigated in many contexts with different alias, such as “anomaly detection” [15], “one-class classification” [18], [19], “open-set recognition” [20], or “novelty detection” [21]. Significant results have been achieved by conventional methods in low-dimensional spaces [19], [22], and some of these methods have also been applied to NLU systems [23], [24].

Some recent neural models use only IND data for OOD detection. Most of these methods follow the threshold-based protocol, and various approaches for calculating the detection scores are devised. Popular approaches include modeling the probability density [8], [25], computing reconstruction losses [7], [26], [27], using classifier ensembles [13], [28], applying Bayesian models [29], relying on distances to nearest-neighbors [9], [30], or even explicitly learning a detection score [31]. Some KNN-based methods are also applied to handle text inputs [32], [33]. However, most of these methods are computationally expensive either in training or inference, and cannot take full advantage of unlabeled data to improve the OOD detection performance. Some of these methods also require a tremendous amount of memories as the number of classes increases [33]. All these disadvantages limit the feasibility of these methods in practical applications.

Another type of neural-based OOD detection model aims to utilize a set of OOD data in the training phase. Specifically, a special “OOD” label is added in a binary or multi-class classifier (e.g., [34]), and the inputs that fall into this special “OOD” class is rejected. However, the feasibility of this naive approach is limited in practice since suitable OOD data are usually hard to

collect, and incorporating too many irrelevant OOD samples in training may cause a serious issue of data imbalance. Further, the OOD distribution is usually too broad to capture with these limited OOD data.

Our study is also related to a large number of works on controllable text generation [35]–[37], some of which also involve an adversarial training process [38], [39] and controls the generated content. However, most previous studies for controllable text generation aim to model a smooth representation space and produce fluent utterances within the data distribution, whereas our model targets at improving the OOD detection performance of an NLU system and tries to generate effective OOD samples that are not in the given data distribution.

Another branch of related studies is the Positive and Unlabeled (PU) learning [40], in which a learning algorithm has only access to positive examples and unlabeled data. The difference between our study and PU learning models is that we aim to reject all samples that are not from IND classes. There is no guarantee that the unlabeled data can cover the entire OOD distribution, which is usually too large to tackle. However, the negative distribution considered in PU learning models is assumed to be completely covered by the unlabeled data.

There are two closely relevant studies from Ryu *et al.* [16] and Lee *et al.* [14]. These studies utilize a GAN based generator to produce OOD samples when building the OOD detector. The major difference between our study and these works is that we focus on generating discrete token sequences, whereas previous works can only generate samples in the continuous space (e.g., images or continuous feature vectors). Moreover, our model can also utilize unlabeled data to improve OOD detection performance, while previous works only use IND data.

III. MODEL

A. Task Definition

In this study, we aim at improving the OOD detection performance of a practical NLU module in a dialogue system. We focus on the intent classification task since it is the most important role of an NLU module. These OOD inputs that are not supported by the current system are rejected once they are detected in the intent classification process.

Our task can be formally defined as below: Given a set of IND data, which are drawn independently from an IND distribution \mathcal{P}_{ind} , i.e.,

$$\mathcal{D}_{ind} = \{(x_1, y_1), \dots, (x_n, y_n)\} \sim \mathcal{P}_{ind},$$

and a set of unlabeled data, which are either drawn from the IND distribution \mathcal{P}_{ind} or the OOD distribution \mathcal{P}_{ood} , i.e.,

$$\mathcal{D}_{mix} = \{\hat{x}_1, \dots, \hat{x}_n\} \sim \mathcal{P}_{ind} \text{ or } \mathcal{P}_{ood},$$

where x_i and \hat{x}_i represent utterances, and $y_i \in \{l_1, \dots, l_m\}$ is x_i 's label (i.e., intent type). We aim to build an intent classifier which can (1) reject the input x if x is drawn from \mathcal{P}_{ood} , and (2) predict the correct intent type of x if it is from \mathcal{P}_{ind} .

Note that in most cases, \mathcal{P}_{ood} is not known, or its underlying space is too large to explore. We cannot expect to capture the

entire \mathcal{P}_{ood} with the limited OOD data sampled from \mathcal{D}_{mix} . However, we can expect to gain some prior knowledge about the OOD distribution with the help of \mathcal{D}_{mix} to improve the performance of OOD detection.

B. Classifier With Entropy Regularization

In this study, the threshold based approach is used for detecting OOD inputs. Similar to the method introduced in [11], we use the maximum value of the Softmax output as the OOD detection score. Specifically, the intent classifier is built with a Softmax output layer to predict an m -dimension distribution $P_\theta(\mathbf{y}|x)$ for each input utterance x :

$$P_\theta(\mathbf{y}|x) = [P_\theta(y = l_1|x), P_\theta(y = l_2|x), \dots, P_\theta(y = l_m|x)]$$

The detection score for x is obtained by:

$$Score(x) = \max_{i \in \{1, 2, \dots, m\}} P_\theta(y = l_i|x) \quad (1)$$

where θ denotes the parameters of the classifier, and m is the number of intent types in the NLU module.

In order to determine whether an input utterance x is from \mathcal{P}_{ind} or \mathcal{P}_{ood} , a threshold t is chosen (usually based on the validation set). The input x is regarded as an IND sample and further processed by the system if $Score(x) \geq t$, otherwise it is determined to be an OOD sample and thus rejected by the system. Therefore, it is desirable for IND inputs to obtain higher detection scores, while OOD inputs to obtain lower detection scores.

Usually, the intent classifier is trained by minimizing the cross-entropy loss:

$$\mathcal{L}_{ce}(\theta) = \mathbb{E}_{(x_i, y_i) \sim \mathcal{P}_{ind}} [-\log P_\theta(y = y_i|x_i)] \quad (2)$$

Minimizing $\mathcal{L}_{ce}(\theta)$ on \mathcal{D}_{ind} enforces the classifier to produce confident predictions on IND samples, which leads to high detection score. However, neural models trained based on the cross-entropy loss tend to be overconfident [41], which results in the fact that samples from \mathcal{P}_{ood} may also receive a high detection score [12] if they share similar patterns and phrases with some IND samples. For example, consider a smartphone without a Bluetooth module. The intelligent assistant equipped on the phone may receive the OOD utterance ‘‘Turn on the Bluetooth please’’ from the user, who may not be familiar with the precise capability of the smartphone. This OOD input may receive a high detection score since the pattern ‘‘Turn on ... please’’ is commonly used in other IND inputs. This makes the IND and OOD inputs indistinguishable with the detection scores. In order to address this issue, a regularization term can be added to enforce a high entropy for the samples from \mathcal{P}_{ood} , i.e.,

$$\mathcal{L}_{ent}(\theta) = \mathbb{E}_{\hat{x} \sim \mathcal{P}_{ood}} [-\mathcal{H}(P_\theta(\mathbf{y}|\hat{x}))] \quad (3)$$

where \mathcal{H} is the Shannon entropy of the predicted distribution. This term is similar to the *confidence loss* used in [14] as it enforces the predicted distribution of OOD inputs closer to the uniform distribution, and thus leads to lower detection scores for OOD inputs.

The total loss for the intent classifier is

$$\mathcal{L}_{cls}(\theta) = \mathcal{L}_{ce}(\theta) + \alpha \mathcal{L}_{ent}(\theta) \quad (4)$$

where α is a hyper-parameter to balance the contribution of the entropy regularization term. In this study, we set $\alpha = 1$.¹

Note that in the original work of [15], $\mathcal{L}_{ent}(\theta)$ is optimized with samples drawn from \mathcal{P}_{ood} . Ideally, we should sample all types of OOD inputs if we cannot obtain any prior knowledge for \mathcal{P}_{ood} . However, this is often infeasible, if not impossible, in practical applications. We thus propose to tackle this issue with a pseudo OOD sample generation module, which will be detailed in the next section.

C. Pseudo OOD Sample Generation

In this section, we present a novel pseudo OOD sample generation (POG) model, which employs an adversarial generation process. The produced pseudo OOD samples can be used to evaluate the entropy regularization term $\mathcal{L}_{ent}(\theta)$, thereby improving the performance of OOD detection.

1) *Model Overview*: The overall architecture of the proposed POG model is shown in Fig. 1. Three major components are included: 1) an autoencoder, 2) an adversarial generation module, and 3) an auxiliary classifier.

The idea of generating effective pseudo OOD samples originates from this observation: most OOD samples look similar to IND inputs (i.e., sharing the same phrases or patterns) but do not correspond to any IND intents. Such OOD samples are usually harder to detect, and more efficient for improving the OOD detection performance if they are used to optimize the ER term [14]. Therefore, in this study, we use the autoencoder to map an input utterance x into a latent code z and use the adversarial generation module to imitate the real latent code z with a fake one \tilde{z} . In this way, the generated utterance \tilde{x}' will look similar to the reconstructed utterance x' . Further, the auxiliary classifier is trained to predict the correct intent label associated with the reconstructed utterance x' and used to regularize the latent code generator to make sure the generated utterance \tilde{x}' will have indistinguishable intent labels and thus being OOD.

Each component is detailed below.

2) *Autoencoder*: This component intends to capture a continuous latent space for the input utterances so that we can apply gradient-based approaches more easily. It contains two functions: An encoder Enc_ϕ (parameterized by ϕ) that maps an input utterance x to a latent code z , i.e.,

$$z = \text{Enc}_\phi(x),$$

and an decoder $P_\psi(x|z)$ (parameterized by ψ) that reconstructs an utterance x' out of z . The autoencoder is trained by minimizing the reconstruction loss:

$$\begin{aligned} \mathcal{L}_{rec}(\phi, \psi) &= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I})} \mathbb{E}_{x \sim \mathcal{P}_{ind}} [-\log P_\psi(x|z + \epsilon)] \\ &= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I})} \mathbb{E}_{x \sim \mathcal{P}_{ind}} [-\log P_\psi(x|\text{Enc}_\phi(x) + \epsilon)] \end{aligned} \quad (5)$$

¹We also tried other values for α , but the experiments show that the final result is not sensitive to this value.

Note that in order to smooth the latent space, we add a Gaussian noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ to z before feeding it to the decoder in the training phase, i.e., the input utterance x is first encoded into a continuous vector z , and then a noise ϵ is sampled from the standard normal distribution $\mathcal{N}(0, \mathbf{I})$. A reconstructed utterance x' is generated based on the latent code $z + \epsilon$ and the autoencoder is optimized to minimize the reconstruction loss shown in Eq. (5). This approach is reported to be effective for smoothing the produced latent space of an autoencoder [38].

3) *Adversarial Generation Module*: This component uses an adversarial training process to approximate the latent codes corresponding to the IND data. Two functions are involved: a generator G_ξ which maps a noise $\epsilon \sim \mathcal{N}$ to a latent code \tilde{z} , i.e.,

$$\tilde{z} = G_\xi(\epsilon),$$

and a discriminator D_η which distinguishes the *real* latent code z from the *generated* latent code \tilde{z} . Intuitively, the generator aims to fool the discriminator, while the discriminator aims to discriminate real codes from generated ones. Once the training is done, the generated latent code \tilde{z} is expected to be able to produce utterances that are similar to the IND inputs.

In this study, we train our generator and discriminator by minimizing the Wasserstein-1 distance [42] between the generated distribution and the data distribution. Specifically, the loss for the generator is

$$\mathcal{L}_g(\xi) = \mathbb{E}_{\epsilon \sim \mathcal{N}} [-D_\eta(G_\xi(\epsilon))] \quad (6)$$

whereas the loss for the discriminator is:

$$\mathcal{L}_d(\eta) = \mathbb{E}_{\epsilon \sim \mathcal{N}} [D_\eta(G_\xi(\epsilon))] - \mathbb{E}_{x \sim \mathcal{P}_{ind}} [D_\eta(\text{Enc}_\phi(x))] \quad (7)$$

In order to enforce the 1-Lipschitz constraint [42] on the discriminator D_η , we employ the gradient penalty term as proposed in [43]:

$$\mathcal{L}_{gp}(\eta) = \mathbb{E}_{\hat{x} \sim \mathcal{P}_{\hat{x}}} [(\|\nabla_{\hat{x}} D_\eta(\hat{x})\|_2 - 1)^2] \quad (8)$$

where the sampling process from $\mathcal{P}_{\hat{x}}$ is approximated by uniformly interpolating between two random samples, one from the data distribution and another from the generated distribution [43].

4) *Auxiliary Classifier*: This component of our model maps a decoded utterance to an m -dimension label distribution, where m is the number of IND intent labels. The parameters ω of the auxiliary classifier (AC) $P_\omega(\mathbf{y}|x')$ are optimized with the cross-entropy loss to predict the correct intent label associated with the utterances x' decoded from the ‘‘real’’ latent code. Specifically, for an IND input x_i (with intent label y_i) that is sampled from the training data, the corresponding ‘‘real’’ latent code $z_i = \text{Enc}_\phi(x_i)$ is first generated using the encoder and a reconstructed utterance x'_i is decoded based on z_i . Then we feed the decoded utterance x'_i to the AC and try to predict its intent label y_i . The following loss is optimized:

$$\mathcal{L}'_{ce}(\omega) = \mathbb{E}_{\substack{(x_i, y_i) \sim \mathcal{P}_{ind}, \\ x'_i \sim P_\psi(x|\text{Enc}_\phi(x_i))}} [-\log P_\omega(y = y_i|x'_i)] \quad (9)$$

Further, we use the AC to guide our latent code generator G_ξ to produce latent codes that can be decoded into OOD samples,

Algorithm 1: Training Procedure of POG.

```

1: for each training iteration do
   /* (1) Train the autoencoder (Encφ, Decψ)
   and the auxiliary classifier (ACω) */
2: Sample {xi, yi}i=1M ~ Dind
3: Compute zi = Encφ(xi), i = 1, ..., M
4: Add Gaussian noise zi ← zi + ε, ε ~ N(0, I)
5: Update φ, ψ by minimizing Lrec(φ, ψ)
6: Decode x'i ~ Pψ(x|zi)
7: Update ω by minimizing L'ce(ω)
   /* (2) Train the discriminator (Dη) */
8: Sample {xi, yi}i=1M ~ Dind, {εi}i=1M ~ N
9: Compute zi = Encφ(xi), and z̃i = Gξ(εi)
10: Update η by minimizing Ld(η) + Lgp(η)
   /* (3) Train the generator (Gξ) */
11: Sample {εi}i=1M ~ N
12: Compute z̃i = Gξ(εi)
13: Update ξ by minimizing Lg(ξ)
14: Decode x̃'i ~ Pψ(x|z̃i)
15: Update ξ by minimizing L'ent(ξ)
16: end for

```

i.e., the generator G_ξ is optimized to enforce a high entropy for the predicted distribution of the AC on the utterance \tilde{x}' decoded from the “fake” latent code. Specifically, we first sample a noise $\epsilon \sim \mathcal{N}$ from the standard normal distribution and generate a “fake” code $\tilde{z} = G_\xi(\epsilon)$ using the generator. An utterance \tilde{x}' is decoded based on \tilde{z} (i.e., $\tilde{x}' \sim P_\psi(x|\tilde{z})$) and the following loss is optimized:

$$\mathcal{L}'_{ent}(\xi) = \mathbb{E}_{\substack{\epsilon \sim \mathcal{N}, \\ \tilde{x}' \sim P_\psi(x|G_\xi(\epsilon))}} [-\mathcal{H}(P_\omega(\mathbf{y}|\tilde{x}'))] \quad (10)$$

where \mathbf{y} denotes the intent type space.

Note that the latent code generator G_ξ trained using the loss $\mathcal{L}_g(\xi)$ and $\mathcal{L}'_{ent}(\xi)$ is trying to accomplish two adversarial targets: **First**, the adversarial loss $\mathcal{L}_g(\xi)$ forces the generated latent code to be close to the IND space, and thus makes the decoded utterance \tilde{x}' looks similar to utterances in \mathcal{D}_{ind} ; **Second**, the regularization loss $\mathcal{L}'_{ent}(\xi)$ ensures the intent associated with the decoded utterance \tilde{x}' cannot be predicted by the AC. Specifically, $\mathcal{L}'_{ent}(\xi)$ reaches its minimum when the AC produces a uniform distribution, which means \tilde{x}' does not belong to any existing intent labels. It is expected that these losses can guide our model to generate OOD samples near the IND distribution (i.e., look similar to the IND samples), thereby making the model more effective in OOD detection.

The training process of the POG model is detailed in Algorithm 1. Note that the text sample \tilde{x}'_i which is decoded in Step 1 of Algorithm 1 is discrete and non-differentiable, which hinders the gradients back-propagating from the AC to the generator G_ξ through the loss $\mathcal{L}'_{ent}(\xi)$ in Step 1. In order to address this issue, we use a continuous approximation approach to replace the token (i.e., one-hot vector) sampled at each time step in \tilde{x}'_i (in Step 1) and x'_i (in Step 1) with the probability vector produced by

Algorithm 2: Training Procedure of AEPOG.

```

1: Each training iteration additionally:
   /* (1b) Train the autoencoder (Encφ, Decψ) */
2: Sample {x̂i}i=1M ~ Dmix
3: Compute zi = Encφ(x̂i), i = 1, ..., M
4: Add Gaussian noise zi ← zi + ε, ε ~ N(0, I)
5: Update φ, ψ by minimizing Lrec(φ, ψ)
   /* (2b) Train the discriminator (Dη) */
6: Sample {x̂i}i=1M ~ Dmix, {εi}i=1M ~ N
7: Compute zi = Encφ(x̂i), and z̃i = Gξ(εi)
8: Update η by minimizing Ld(η) + Lgp(η)

```

the decoder $P_\psi(x|z)$. These “soft” tokens are fed into the AC to make the whole computation process differentiable. Specifically, the word embedding fed at each time step of the AC is computed as an average over all the word embeddings weighted by the input probability distribution. Moreover, we also applied the temperature scaling technique to sharp the output distribution, i.e., the logits in each time step are divided by a temperature of t to produce the output distribution. In our experiments, the value of t decreases from 1 to 0 as the training proceeds. A small value of t sharpens the output distribution to make it close to a one-hot vector.

D. Utilizing Unlabeled Data

It is observed in previous studies that if the OOD data used in the training and testing phase are similar, the OOD detection performance is better [14], [15]. However, using human-labeled OOD data in Eq. (3) is expensive, and it is hard for the POG model to gain any prior knowledge about the testing OOD samples if only IND data is utilized. However, fortunately, the unlabeled data \mathcal{D}_{mix} , which are usually easier to collect, can help us to solve this issue.

In this study, we augment the training process of the autoencoder in POG with the data from \mathcal{D}_{mix} . It is expected that the latent space produced by the autoencoder will be enriched with OOD samples, and the AC will lead the generator to focus on these OOD samples in the adversarial training process. Specifically, this involves adding a few additional steps in Algorithm 1, in particular, two steps are added in each training iteration: (1b) update ϕ, ψ with the reconstruction loss and (2b) update η with the discriminator loss utilizing the data sampled from \mathcal{D}_{mix} . We denote the augmented model as AEPOG, and the corresponding training algorithm is shown in Algorithm 2. Experiments in Section IV demonstrate that the generated pseudo OOD utterances can be more effective at improving the OOD detection performance.

IV. EXPERIMENT**A. Dataset**

We evaluated the proposed model on two datasets:

- 1) *OSQ dataset* [44]: this dataset covers 150 IND intents and also provides a set of manually labeled Out-of-Scope

TABLE I
STATISTICS OF THE OSQ AND IPA DATASETS. \mathcal{D}_{mix} DENOTES A MIXTURE OF
IND AND OOD DATA

		Train	Validate	Test
OSQ Dataset	IND	15.00K	3.00K	4.50K
	OOD	-	0.10K	1.00K
	\mathcal{D}_{mix}	10.25K	-	-
IPA Dataset	IND	28.90K	3.60K	3.60K
	OOD	-	1.20K	1.20K
	\mathcal{D}_{mix}	20.00K	-	-

Queries (OSQ) that are not supported by the current system. These out-of-scope queries are regarded as OOD data in our experiments.

- 2) *IPA dataset*: this dataset contains a set of Chinese utterances that were collected and annotated in the development process of a commercialized Intelligent Personal Assistant (IPA) named Bixby.² This dataset covers 67 domains, which can be further divided into 1,310 intents. In this study, 310 intents are randomly selected as the OOD intents, and the corresponding utterances are used as OOD data. Domain level labels are used for the IND utterances, i.e., 67 class labels are used for these IND data, and the utterances for each IND intent is assigned with the corresponding domain label.

The IND data were divided into the train, validation, and test sets, while the OOD data were only used for validation and test. In addition to the labeled data, some unlabeled data (i.e., a mixture of IND and OOD data) \mathcal{D}_{mix} were also collected for the AEPOG model. Specifically, for the OSQ dataset, a mixture of 10 K IND data and 250 OOD data were used as \mathcal{D}_{mix} . For the IPA dataset, 20 K unlabeled utterances were extracted from user logs and used as \mathcal{D}_{mix} . The partitions of the data used in our experiment are shown in Table I.

Note that the data partition scheme for the IPA dataset to define OOD inputs ensures that the OOD and IND utterances have similar patterns. This simulates the actual situations that we have met in our real system, i.e., users may issue OOD utterances that are similar to IND utterances since they are not familiar with the precise capabilities of the system. This data partition scheme makes the OOD detection task more difficult, and it demonstrates a lower bound for the OOD detection performance of a real NLU system. Measuring and guaranteeing such lower bounds are essential for practical systems.

In order to better demonstrate the performance of our method on the IPA dataset, we collect an extra set of test data, i.e., Chat dataset, for the models trained on the IPA dataset. This dataset covers a broad range of OOD utterances such as chitchat or nonsense, which are commonly received by a practical NLU system. The Chat dataset contains 1.2 K utterances, and it is only used in the test phase of our experiment. The utterances in the Chat dataset are manually annotated based on the user logs.

²Bixby is shipped with a large number of Samsung products, such as phones, refrigerators, or watches.

B. Model Implementation

Our text classifier was implemented using the CNN architecture [45]. Four kernel sizes (2, 3, 4, 5) were used, and each kernel had 128 feature maps. A three-layer MLP was added on top of the CNN feature, and the hidden size of each layer was 512. The vocabulary size was of 8.3 K and 18.2 K for the classifiers on the OSQ and IPA datasets, respectively. Note that the vocabulary we used covers all the tokens in the IND and OOD data. Pre-trained word vectors with the dimension of 100 were used to initialize the word embeddings. Specifically, 81.46% and 70.01% percent of word embeddings were initialized with the pre-trained word vectors for the IND and OOD datasets, respectively. The rest word embeddings were randomly initialized. All the classifier was trained for 30 epochs.

For the POG model, the auxiliary classifier shared the same structure and hyper-parameter setting with the text classifier. The encoder and decoder were both implemented using LSTM [46]. The hidden size of the LSTM used for the OSQ and IPA datasets was 100 and 256, respectively. The generator and discriminator were both four-layer MLPs activated with the Leaky ReLU function [47]. The MLP hidden size was 512 and 1024, and the vocabulary size was of 5.8 K and 14.1 K, respectively, for the OSQ and IPA datasets. It covers all the tokens in the IND training data. The POG and AEPOG models were trained for 80 epochs.

C. Baselines

In this work, we used several threshold-based OOD detectors as baselines. These baselines differ mainly in the way to calculate detection scores.

- 1) *MSP [11]*: A text classifier with a Softmax output layer is trained, and the maximum Softmax output is used as the detection score (i.e., calculated using Eq. (1)). The training objective of this classifier does not include the entropy regularization (ER) term.
- 2) *Entropy*: The Shannon entropy of the predicted distribution for each input is used as the detection score. Higher entropy means higher uncertainty of the prediction, which indicates that the input sample may be from the OOD data.
- 3) *ODIN [12]*: The temperature scaling and input perturbation technique is applied to the text classifier as obtained in the previous **MSP** baseline, and the maximum Softmax output is used as the detection score. Note that small perturbations of each input are added to the last feature layer in this baseline.
- 4) *DOC [28]*: m binary classifiers are built for m classes. The maximum confidence score predicted by these m classifiers is used as the detection score.
- 5) *KNN*: The feature for each input sample is extracted using a pre-trained classifier (i.e., the classifier as obtained in the previous **MSP** baseline), and the Euclid distance of each feature vector to its nearest class is used as the detection score.
- 6) *Maha. Dis. [48]*: The Mahalanobis distance is used to replace the Euclid distance in the previous **KNN** baseline.

- 7) *Cont. GAN* [16]: A GAN-based model is trained to generate continuous OOD features to mimic features extracted from IND samples. The confidence of the discriminator is used as the detection score. A low confidence score suggests that the input sample may be from the OOD data.
- 8) *AE*: An autoencoder is trained on the IND data, and the reconstruction error of each input is used as the detection score.
- 9) *Likelihood Ratio* [9]: The likelihood ratio is used as the detection score. The background model is trained using perturbed IND samples, which are generated by randomly replacing tokens in the IND samples with a probability of 0.5.³
- 10) *ER+Perturb*: An ER term $\mathcal{L}_{ent}(\theta)$ (Eq. (3)) is added to the training objective of the classifier, and this term is optimized using perturbed IND samples. These perturbed samples are constructed in a similar way as in the previous **Likelihood Ratio** baseline.
- 11) *ER+Mix*: The ER term in the training objective is optimized using the data from \mathcal{D}_{mix} . This baseline demonstrates a naive way to utilize \mathcal{D}_{mix} .

Our baselines cover a variety of competitive OOD detection models that are currently available. Note that the text classifier involved in each baseline follows the same structure and hyperparameter setting with our classification model (see Section IV-B), and the baseline **AE** and **Likelihood Ratio** share the same LSTM structure with our POG model.

Several ablation tests are also performed to validate the effect of each component in our model:

- 1) *w/o Noise*: The autoencoder is trained without adding the Gaussian noise in Eq. (5).
- 2) *w/o Soft Token*: The loss $\mathcal{L}'_{ent}(\xi)$ in Eq. (10) is optimized without using the soft token approximation approach. We instead use the policy gradient algorithm to estimate the gradients through the sampling process of the decoder.

Note that all the ablation studies are performed without using the unlabeled data \mathcal{D}_{mix} .

D. Metrics

For OOD detection, we used three common metrics [9], [11], [14], [15]:

- 1) *AUROC*: The area under the receiver operating characteristic (ROC) curve; Higher values are better.
- 2) *AUPR*: The area under the precision-recall curve when OOD inputs are treated as the positive samples; Higher values are better.
- 3) *FPRN*: The false-positive rate (FPR) when the true positive rate (TPR) is $N\%$.⁴ In this study, FPR95 and FPR90 are reported, since such settings are commonly used in practical systems. Lower values are better.

³Other probability values were also tested, but similar results were obtained

⁴Note that $FPR = \frac{FP}{FP+TN}$, and $TPR = \frac{TP}{TP+FN}$, where FP, TP, TN, FN is the number of false positives, true positives, true negatives, false negatives, respectively.

TABLE II

OOD DETECTION PERFORMANCE ON THE OSQ DATASET. EACH RESULT IN THIS TABLE IS AN AVERAGE OF TEN DIFFERENT RUNS. THE NOTATION \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow MEANS LOWER VALUES ARE BETTER. THE MODEL ER+AEPOG IS SIGNIFICANTLY BETTER THAN OTHER MODEL WITH p -VALUE < 0.01 (\dagger) AND p -VALUE < 0.05 ($*$) USING T-TEST

Model	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow	FPR90 \downarrow
Cont. GAN	52.22 \dagger	82.79 \dagger	94.40 \dagger	88.17 \dagger
Maha. Dis.	67.14 \dagger	90.56 \dagger	91.94 \dagger	83.30 \dagger
Likelihood Ratio	85.60 \dagger	96.07 \dagger	62.50 \dagger	43.40 \dagger
AE	87.78 \dagger	96.98 \dagger	58.50 \dagger	40.10 \dagger
MSP	92.86 \dagger	98.24 \dagger	39.72 \dagger	21.76 \dagger
Entropy	92.82 \dagger	98.87 \dagger	31.91 \dagger	19.64 \dagger
KNN	93.33 \dagger	98.20 \dagger	33.86 \dagger	18.78 \dagger
DOC	94.24 \dagger	98.55 \dagger	30.02 \dagger	14.94 \dagger
ODIN	95.14 \dagger	98.84 $*$	26.04 $*$	11.70 \dagger
ER+Perturb	94.01 \dagger	98.55 \dagger	34.04 \dagger	15.32 \dagger
ER+Mix	93.48 \dagger	98.31 \dagger	33.16 \dagger	17.86 \dagger
ER+POG	95.41 $*$	98.94 $*$	25.00 $*$	10.10 \dagger
ER+AEPOG	95.83	99.05	23.70	9.50
w.o. Noise	94.03 $*$	98.53 $*$	35.22 \dagger	17.74 \dagger
w.o. Soft Token	93.85 \dagger	98.53 $*$	39.24 \dagger	17.88 \dagger

Note that the metric FPRN is of more practical value in real-world applications since it evaluates the performance of an OOD detection module at a particular threshold. It is directly related to the performance of a deployed system. Lower FPRN means triggering fewer false alarms when the performance on the IND data is guaranteed with a precision of $N\%$ [15]. On the contrary, the metrics AUPR and AUROC evaluate the performance across various thresholds. In other words, AUPR and AUROC are threshold independent.

E. Effects of Generated Pseudo OOD Utterances

First of all, we evaluated the effectiveness of the pseudo OOD utterances generated by our model POG when only IND data is available. More concretely, the POG model was trained and, meanwhile, a set of pseudo OOD utterances (with the same size of the IND training data) was sampled to optimize the ER term $\mathcal{L}_{ent}(\theta)$ (i.e., Eq. (3)). The performance of the resulting model (i.e., ER+POG) on the OSQ and IPA datasets is depicted in Table II and Table III, respectively. Table IV shows the test results on the Chat dataset for the models trained on the IPA dataset.

It can be seen that our model outperforms all the baselines on both datasets significantly, which demonstrates the effectiveness of the generated pseudo OOD utterances. It is also interesting to see that: **1)** Detecting OOD samples on the IPA dataset is more difficult compared to the OSQ dataset (see Table II and Table III). It is because that most OOD inputs in the IPA test data look similar to the IND inputs. Moreover, the models trained on the IPA dataset obtain a large performance gain when testing on the Chat dataset, especially on the FPRN metric. This validates our claim that the OOD utterances that look similar to IND utterances are more difficult to detect. Further, the fact that the proposed model ER+POG surpasses all the baselines on both datasets also proves its ability to handle various kinds of inputs with different

TABLE III

OOD DETECTION PERFORMANCE ON THE IPA DATASET. EACH RESULT IN THIS TABLE IS AN AVERAGE OF TEN DIFFERENT RUNS. THE NOTATION \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow MEANS LOWER VALUES ARE BETTER. THE MODEL ER+AEPOG IS SIGNIFICANTLY BETTER THAN OTHER MODEL WITH p -VALUE < 0.01 (MARKED BY \dagger) AND p -VALUE < 0.05 (*) USING T-TEST

Model	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow	FPR90 \downarrow
Cont. GAN	58.06 \dagger	80.47 \dagger	91.51 \dagger	85.05 \dagger
Maha. Dis.	45.65 \dagger	23.45 \dagger	95.01 \dagger	89.97 \dagger
Likelihood Ratio	69.65 \dagger	86.21 \dagger	84.44 \dagger	74.46 \dagger
AE	67.57 \dagger	85.91 \dagger	86.77 \dagger	74.79 \dagger
MSP	72.66 \dagger	86.42 \dagger	77.79 \dagger	86.42 \dagger
Entropy	72.87 \dagger	86.42 \dagger	88.15 \dagger	75.41 \dagger
KNN	67.94 \dagger	82.62 \dagger	76.87 \dagger	63.43 \dagger
DOC	71.68 \dagger	46.03 \dagger	79.40 \dagger	64.34 \dagger
ODIN	72.90 \dagger	86.53 \dagger	77.57 \dagger	63.28 \dagger
ER+Perturb	73.24 \dagger	86.89*	77.88 \dagger	63.30 \dagger
ER+Mix	33.71 \dagger	63.55 \dagger	96.44 \dagger	92.33 \dagger
ER+POG	73.83 \dagger	87.15*	76.17 \dagger	62.28 \dagger
ER+AEPOG	75.86	87.95	71.67	56.78
w/o Noise	71.01*	85.10*	79.17 \dagger	65.17 \dagger
w/o Soft Token	72.09*	86.31*	80.22 \dagger	65.07 \dagger

TABLE IV

OOD DETECTION PERFORMANCE ON THE IPA DATASET WHEN TESTING WITH THE CHAT DATASET. EACH RESULT IN THIS TABLE IS AN AVERAGE OF TEN DIFFERENT RUNS. THE NOTATION \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow MEANS LOWER VALUES ARE BETTER. THE MODEL ER+AEPOG IS SIGNIFICANTLY BETTER THAN OTHER MODEL WITH p -VALUE < 0.01 (MARKED BY \dagger) AND p -VALUE < 0.05 (*) USING T-TEST

Model	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow	FPR90 \downarrow
Cont. GAN	69.80 \dagger	87.55 \dagger	90.77 \dagger	76.18 \dagger
Maha. Dis.	73.85 \dagger	48.20 \dagger	73.96 \dagger	61.48 \dagger
Likelihood Ratio	90.83 \dagger	96.63 \dagger	41.51 \dagger	27.37 \dagger
AE	92.41 \dagger	97.21 \dagger	37.77 \dagger	18.97 \dagger
MSP	88.41 \dagger	95.05 \dagger	52.30 \dagger	32.56 \dagger
Entropy	89.02 \dagger	95.05 \dagger	37.56 \dagger	29.31 \dagger
KNN	89.43 \dagger	95.02 \dagger	36.17 \dagger	22.91 \dagger
DOC	93.03 \dagger	97.28 \dagger	35.44 \dagger	18.07 \dagger
ODIN	90.25 \dagger	95.67 \dagger	37.92 \dagger	24.29 \dagger
ER+Perturb	96.28 \dagger	98.72*	19.23 \dagger	10.80 \dagger
ER+Mix	92.76 \dagger	97.27 \dagger	38.44 \dagger	21.71 \dagger
ER+POG	98.53*	99.42*	7.70*	4.36*
ER+AEPOG	98.60	99.46	7.64	4.24
w/o Noise	94.26 \dagger	98.07 \dagger	33.08 \dagger	18.56 \dagger
w/o Soft Token	94.08 \dagger	98.06 \dagger	40.07 \dagger	16.96 \dagger

distributions. **2)** The results for the ablation studies validate the effect of the Gaussian noise and soft token approximation approach used in our study. A significant performance drop is observed on both datasets. **3)** The proposed ER+POG model obtains large improvements on the FPR N metrics. In particular, the relative performance gain for FPR95 and FPR90 on the OSQ dataset reaches 4.16% and 15.84%, respectively, compared to the best performing baselines. Even on the more difficult IPA dataset, the relative performance gain also reaches 0.92% and 1.61% for FPR95 and FPR90, respectively. This indicates that the proposed model is more suitable in real-world applications since the metric FPR N directly reflects the performance of the

TABLE V

CLASSIFICATION ACCURACY ON IND INPUTS FOR EACH MODEL. EACH RESULT IN THIS TABLE IS AN AVERAGE OF TEN DIFFERENT RUNS. THE MODEL ER+AEPOG IS SIGNIFICANTLY BETTER THAN MOST OF OTHER MODELS WITH p -VALUE < 0.01 (MARKED BY \dagger) AND p -VALUE < 0.05 (*) USING T-TEST

Model	OSQ Dataset	IPA Dataset
Cont. GAN	70.43 \dagger	84.68 \dagger
Maha. Dis.	82.36 \dagger	89.73 \dagger
MSP	92.61*	91.19
KNN	81.05 \dagger	88.32 \dagger
DOC	91.84 \dagger	89.20 \dagger
ER+Perturb	90.96 \dagger	91.22
ER+Mix	89.77 \dagger	90.41 \dagger
ER+POG	93.31	91.37
ER+AEPOG	93.32	91.40
w/o Noise	93.01	91.34
w/o Soft Token	93.10	90.96*

deployed models. **4)** The Cont. GAN baseline proposed in [16] only performs slightly better than random guesses. This shows that the confidence score of the discriminator is not efficient in detecting OOD inputs, and using GAN to generate continuous features is not helpful to improve the OOD detection performance in this baseline.

It is also worth mentioning that the performance of the NLU model on the IND input classification task is not affected if we optimize the ER term with utterances generated by the proposed model (see Table V). Notably, in some cases, the utterances generated by the proposed model even help to improve the classification accuracy of the NLU module on IND inputs. This may be because optimizing the ER term with utterances produced by the POG or AEPOG model helps to prevent the classifier from over-fitting, and it further demonstrates the effectiveness of the proposed model in practical applications.

F. Effects of Utilizing Unlabeled Data

We also verified the effectiveness of our model when a set of unlabeled data \mathcal{D}_{mix} is available. Specifically, an AEPOG model was trained by optimizing $\mathcal{L}_{\text{rec}}(\phi, \psi)$ with the data from \mathcal{D}_{mix} (Algorithm 2), and the sampled pseudo OOD utterances were used to optimize $\mathcal{L}_{\text{ent}}(\theta)$ when training the intent classifier. The performance of the resulting classifier on the OSQ and IPA datasets is shown in the last line of Table II and Table III, respectively.

It can be seen that the model ER+AEPOG surpasses all the other models significantly, especially on the FPR N metric. Particularly, on the OSQ dataset, the performance of the ER+AEPOG model improves 9.87% and 23.16% (relatively) compared to the best performing baselines on FPR95 and FPR90, respectively. On the IPA dataset, the performance improves 7.26% and 11.45% (relatively) for FPR95 and FPR90, respectively. This verifies our claim that the proposed AEPOG model can effectively utilize unlabeled data to improve the effectiveness of the generated OOD samples. Thus it is more suitable in practical applications.

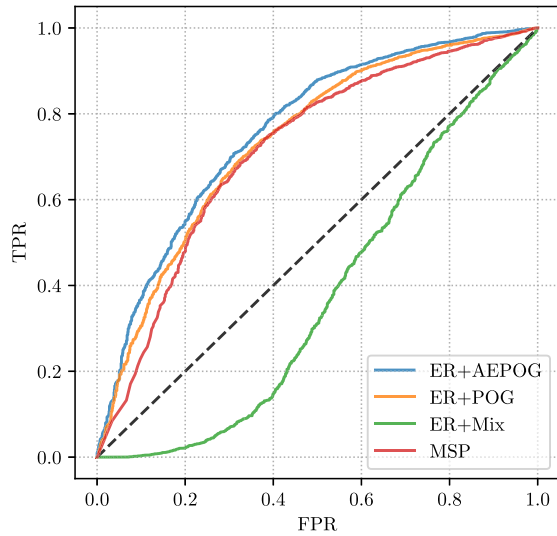


Fig. 2. ROC curves for the models on the IPA dataset. The model ER+AEPOG improves the AUROC of OOD detection compared to ER+POG, whereas a sharp decrease for the AUROC of OOD detection is observed for the model ER+Mix, which directly utilize \mathcal{D}_{mix} to optimize the ER term of the classifier.

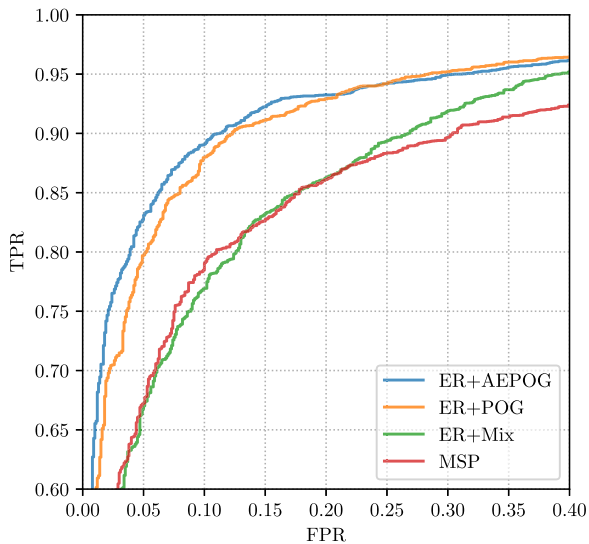


Fig. 3. ROC curves for the models on the OSQ dataset. We zooms in the upper-left corner of the ROC curves to facilitate a clearer view.

It is also interesting to see that the baseline model ER+Mix, which optimizes the ER term $\mathcal{L}_{\text{ent}}(\theta)$ directly using the data from \mathcal{D}_{mix} when training the classifier, experiences a remarkable performance drop on all metrics, especially on the IPA dataset (see Fig. 2). This indicates that the ER term itself cannot utilize \mathcal{D}_{mix} directly. It further demonstrates that the performance improvement obtained by the ER+AEPOG model attributes to the proposed AEPOG model since the proposed adversarial training process helps to produce more effective pseudo OOD samples to improve the OOD detection performance.

Fig. 3 shows the ROC curves obtained from these models on the OSQ dataset. The proposed ER+POG model surpasses all the baselines, and ER+AEPOG further improves the performance

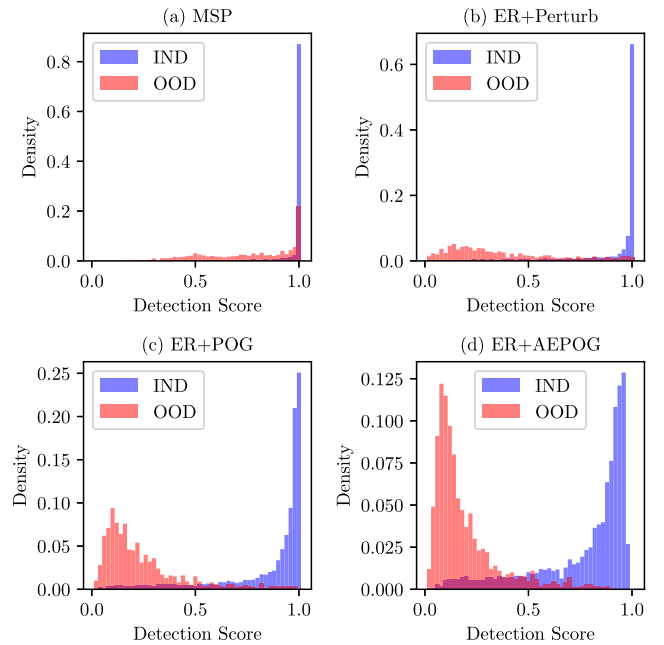


Fig. 4. Distributions of detection scores corresponding to the IND and OOD samples of the OSQ dataset.

of OOD detection. Similar results are also obtained on the IPA dataset (see Fig. 2). Note that the performance improvement of the ER+AEPOG model over the ER+POG model is relatively small on the OSQ dataset compared to the IPA dataset. This is partly because the OOD data only takes a small proportion (i.e., 2.44%) of \mathcal{D}_{mix} in the OSQ dataset, whereas about one third (33.33%) intents are randomly selected as OOD in the IPA dataset.

We also analyzed the distribution of the detection scores obtained from different models on the OSQ and IPA datasets. Specifically, for the MSP model, the distributions corresponding to IND and OOD inputs are closely overlapped (see Fig. 4(a) and Fig. 5(a)). Marginal improvements are observed for the model ER+Perturb, which optimizes the ER term using randomly perturbed IND inputs (Fig. 4(b) and Fig. 5(b)). This indicates that the ER term helps to improve the OOD detection performance, but only to a limited extent. Better separation of the detection scores between the IND and OOD inputs is observed when the ER term is optimized with the pseudo OOD utterances generated by the POG model (Fig. 4(c) and Fig. 5(c)), and this separation is enlarged if the utterances generated by the AEPOG model are used (Fig. 4(d) and Fig. 5(d)). This indicates that the pseudo OOD utterances generated by the proposed model facilitate OOD detection, and the effectiveness of these pseudo OOD utterances can be further improved when utilizing the unlabeled data.

G. Feature Vectors for IND and OOD Inputs

To further investigate the benefit of the proposed model on OOD detection, we visualized the feature spaces of different intent classifiers on the OSQ dataset (see Fig. 6). Specifically,

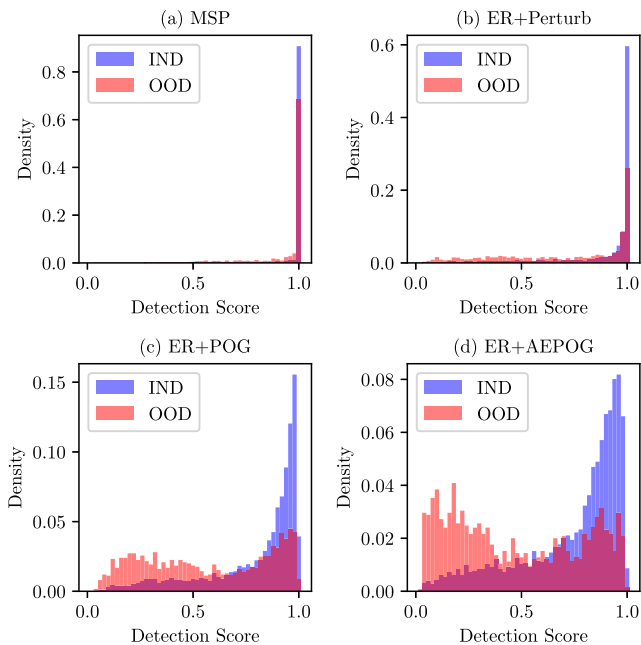


Fig. 5. Distributions of detection scores corresponding to the IND and OOD samples of the IPA dataset.

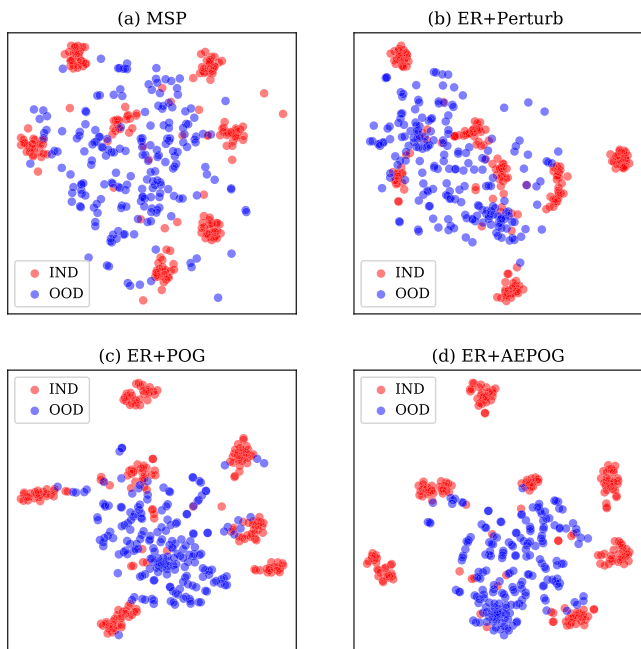


Fig. 6. t-SNE visualization of the features vectors associated with the test samples from the OSQ dataset.

we fed the test samples of the OSQ dataset to each intent classifier and obtained the feature vector of each sample from the penultimate layer of the network. The t-SNE algorithm [49] was used to map these vectors into 2-dimensions.

It can be seen that all the intent classifiers cluster features of the IND samples into separated groups. This coincides with the observation that all models obtain high classification performance on IND inputs (Table V). However, as to the OOD

TABLE VI
UTTERANCES SAMPLED FROM THE IND AND OOD TEST SET OF THE OSQ DATASET (I.E., HUMAN GENERATED UTTERANCES), AND THE PSEUDO OOD UTTERANCES GENERATED USING THE POG AND AEPOG MODEL

IND Samples	Translate hello in French. Locate my phone please. Schedule a gas bill payment. Help me change my insurance plan. I'd like to improve my credit score.
OOD Samples	How much is my car worth used. Can you add a bag to my reservation. How do you fix a leaking sink. How long do wire transfer take. When was Toyota created.
Generated by POG	Please obtain French. How can make my phone get. What meetings can I schedule there. Can you help me an setting. I'd like to include the email my my dinner.
Generated by AEPOG	How much is this payments please. How good is my reservation like. How do you divided for pork. Tell me how long I taken for chilis off. When was today's name with delta vehicle.

samples, the MSP model scatters the features of OOD samples around the features of IND samples (Fig. 6(a)), making it hard for OOD detection. Whereas, the classifiers that are regularized by the ER term produce more distinguishable features for IND and OOD samples (Fig. 6(b)). The utterances generated by the proposed POG and AEPOG model enhance this separation (Fig. 6(c) and Fig. 6(d)). This facilitates the detection of OOD inputs.

H. Case Study for Generated OOD Utterances

Table VI shows some randomly selected cases of the pseudo OOD utterances generated by the proposed POG and AEPOG model on the OSQ dataset. It is interesting to see that the POG model can combine words and phrases from the IND samples (such as “French” or “my phone”) in a grammatical way. This makes the produced pseudo OOD utterances look similar to the IND data but possess indistinguishable intents. Moreover, the AEPOG model can make use of phrases from the OOD data, and thus produces more effective utterances that can be used to improve the OOD detection performance.

V. CONCLUSION

In this paper, we propose a novel model POG to generate pseudo OOD samples that can be used to improve the performance of OOD detection in an NLU module. An autoencoder is used to map each input utterance to a latent code, and an adversarial training process is employed to make the codes of OOD samples similar to those of IND samples. This makes the generated OOD samples look similar to IND samples. Further, an auxiliary classifier is introduced to regularize the generator and thereby ensures these generated pseudo OOD samples to have indistinguishable intents. Our experiments show that the pseudo OOD samples generated by the proposed POG model

can be used to effectively improve the OOD detection performance of the NLU module by optimizing the ER term when training the NLU module. Further, it is also demonstrated that augmenting the training process of the autoencoder in the POG model improves the effectiveness of the generated pseudo OOD samples.

As future works, we will explore this idea in classification problems for longer text since it is more challenging to generate longer documents for OOD detection. Also, note that our study focuses on detecting OOD inputs for the NLU module. It is worth exploring to equip other components of the task-oriented dialogue system, such as dialogue state tracking or dialogue management modules, with the ability to detect OOD inputs.

ACKNOWLEDGMENT

The first author would also like to thank Wonkwang Shin, Seonghan Ryu and Kyoung-Gu Woo for fruitful discussions, and thank the support provided by the Bixby development team.

REFERENCES

- [1] Y.-B. Kim, D. Kim, A. Kumar, and R. Sarikaya, "Efficient large-scale neural domain classification with personalized attention," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 2214–2224.
- [2] R. Sarikaya, "The technology behind personal digital assistants: An overview of the system architecture and key components," *IEEE Signal Process. Mag.*, vol. 34, no. 1, pp. 67–81, Jan. 2017.
- [3] K. M. Yoo, Y. Shin, and S.-G. Lee, "Data augmentation for spoken language understanding via joint variational generation," 2018, in AACL. [Online]. Available: <https://www.aaii.org/ojs/index.php/AAAI/article/view/4729>
- [4] G. Fei and B. Liu, "Breaking the closed world assumption in text classification," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Human Lang. Technol.*, Jun. 2016, pp. 506–514.
- [5] W. J. Scheirer, A. de Rezende Rocha, A. Sapkota, and T. E. Boult, "Toward open set recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 7, pp. 1757–1772, 2013.
- [6] A. Bendale and T. E. Boult, "Towards open set deep networks," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2016, pp. 1563–1572.
- [7] S. Ryu, S. Kim, J. Choi, H. Yu, and G. G. Lee, "Neural sentence embedding using only in-domain sentences for out-of-domain sentence detection in dialog systems," *Pattern Recognit. Lett.*, vol. 88, pp. 26–32, 2017.
- [8] E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan, "Do deep generative models know what they don't know?" in *Proc. Int. Conf. Learn. Representations*, 2019.
- [9] J. Ren, P. J. Liu, E. Fertig, J. Snoek, R. Poplin, "Likelihood ratios for out-of-distribution detection," in *Proc. Adv. Neural Inf. Process. Syst.* 2019. [Online]. Available: <http://papers.nips.cc/paper/9611-likelihood-ratios-for-out-of-distribution-detection>
- [10] A. Shafaei, M. Schmidt, and J. J. Little, "A Less Biased Evaluation of Out-of-distribution Sample Detectors," in *Proc. British Machine Vision Conference*, 2019. [Online]. Available: <https://arxiv.org/abs/1809.04729>
- [11] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [12] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [13] A. Vyas, N. Jammalamadaka, X. Zhu, D. Das, B. Kaul, and T. L. Willke, "Out-of-distribution detection using an ensemble of self supervised leave-out classifiers," in *Proc. Eur. Conf. Comput. Vision*, 2018, pp. 550–564.
- [14] K. Lee, H. Lee, K. Lee, and J. Shin, "Training confidence-calibrated classifiers for detecting out-of-distribution samples," in *Proc. Int. Conf. Learn. Representations*, 2018.
- [15] D. Hendrycks, M. Mazeika, and T. Dietterich, "Deep anomaly detection with outlier exposure," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [16] S. Ryu, S. Koo, H. Yu, and G. G. Lee, "Out-of-domain detection based on generative adversarial network," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2018, pp. 714–718.
- [17] I. J. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [18] S. S. Khan and M. G. Madden, "A survey of recent trends in one class classification," in *Proc. Irish Conf. Artif. Intell. Cogn. Sci.*, 2009, pp. 188–197.
- [19] S. S. Khan and M. G. Madden, "One-class classification: Taxonomy of study and review of techniques," *Knowl. Eng. Rev.*, vol. 29, no. 3, pp. 345–374, 2014.
- [20] C. Geng, S.-j. Huang, and S. Chen, "Recent advances in open set recognition: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9040673>
- [21] M. Kliger and S. Fleishman, "Novelty detection with gan," 2018. [Online]. Available: <https://arxiv.org/abs/1802.10560>
- [22] M. A. F. Pimentel, D. A. Clifton, C. Lei, and L. Tarassenko, "A review of novelty detection," *Signal Process.*, vol. 99, no. 6, pp. 215–249, 2014.
- [23] I. Lane, T. Kawahara, T. Matsui, and S. Nakamura, "Out-of-domain utterance detection using classification confidences of multiple topics," *IEEE Trans. Audio Speech Lang. Process.*, vol. 15, no. 1, pp. 150–161, 2006.
- [24] G. Tur, A. Deoras, and D. Hakkani-Tr, "Detecting out-of-domain utterances addressed to a virtual personal assistant," in *Proc. Interspeech*, Sep. 2014, pp. 283–287.
- [25] S. Pidhorskyi, R. Almoheisen, and G. Doretto, "Generative probabilistic novelty detection with adversarial autoencoders," in *Proc. Advances Neural Inf. Process. Syst.*, 2018, pp. 6822–6833.
- [26] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Special Lecture IE*, vol. 2, pp. 1–18, 2015.
- [27] I. Golan and R. El-Yaniv, "Deep anomaly detection using geometric transformations," in *Proc. Advances Neural Inf. Process. Syst.*, 2018, pp. 9758–9769.
- [28] L. Shu, H. Xu, and B. Liu, "DOC: Deep open classification of text documents," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2017, pp. 2911–2916.
- [29] A. Malinin and M. Gales, "Predictive uncertainty estimation via prior networks," in *Proc. Advances in Neural Inf. Process. Syst.*, 2018, pp. 7047–7058.
- [30] X. Gu, L. Akoglu, and A. Rinaldo, "Statistical analysis of nearest neighbor methods for anomaly detection," in *proc. Neural Information Processing Systems*, 2019. [Online]. Available: <https://papers.nips.cc/paper/9274-statistical-analysis-of-nearest-neighbor-methods-for-anomaly-detection.pdf>
- [31] T. DeVries and G. W. Taylor. [Online]. Available: <https://arxiv.org/abs/1802.04865>
- [32] K. J. Oh *et al.*, "Out-of-domain detection method based on sentence distance for dialogue systems," in *Proc. IEEE Int. Conf. Big Data Smart Comput.*, 2018, pp. 673–676.
- [33] H. Xu, S. L. Liu, Bing, and P. Yu, "Open-world learning and application to product classification," in *Proc. Conf. World Wide Web*, 2019, pp. 3413–3419.
- [34] J.-K. Kim and Y.-B. Kim, "Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisfying false acceptance rates," in *Proc. Interspeech*, 2018, pp. 556–560.
- [35] S. R. Bowman, L. Vilnis, O. Vinyals, A. Dai, R. Jozefowicz, and S. Bengio, "Generating sentences from a continuous space," in *Proc. 20th SIGNLL Conf. Computat. Natural Lang. Learn.*, 2016, pp. 10–21.
- [36] H. Bahuleyan, L. Mou, H. Zhou, and O. Vechtomova, "Stochastic Wasserstein autoencoder for probabilistic sentence generation," in *Proc. North Amer. Chapter Assoc. Comput. Linguistics: Human Lang. Technologies, Volume 1 (Long and Short Papers)*, Jun. 2019, pp. 4068–4076.
- [37] J. Zhao, Y. Kim, K. Zhang, A. M. Alexander, Y. LeCun, "Adversarially regularized autoencoders," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018. [Online]. Available: <http://proceedings.mlr.press/v80/zhao18b.html>
- [38] S. Subramanian, S. R. Mudumba, A. Sordani, A. Trischler, A. C. Courville, and C. Pal, "Towards text generation with adversarially learned neural outlines," in *Proc. Advances Neural Inf. Process. Syst.*, 2018, pp. 7551–7563.
- [39] Z. Hu, Z. Yang, X. Liang, R. Salakhutdinov, and E. P. Xing, "Toward controlled generation of text," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, vol. 70, pp. 1587–1596.
- [40] J. Bekker and J. Davis, "Learning from positive and unlabeled data: A survey," *Machine Learning*, 2020. [Online]. Available: <https://doi.org/10.1007/s10994-020-05877-5>

- [41] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, vol. 70, pp. 1321–1330.
- [42] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 214–223.
- [43] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," in *Proc. Advances Neural Inf. Process. Syst.*, 2017, pp. 5767–5777.
- [44] S. Larson *et al.*, "An evaluation for intent classification and out-of-scope prediction," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, Nov. 2019, pp. 1311–1316.
- [45] Y. Kim, "Convolutional neural networks for sentence classification," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2014. [Online]. Available: <https://www.aclweb.org/anthology/D14-1181/>
- [46] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [47] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. Int. Conf. Mach. Learn.*, 2013.
- [48] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *Proc. Advances Neural Inf. Process. Syst.*, 2018, pp. 7167–7177.
- [49] L. v. d. Maaten and G. Hinton, "Visualizing data using T-SNE," *J. Mach. Learn. Res.*, vol. 9, pp. 2579–2605, 2008.



Yinhe Zheng received the Ph.D. degree from the China University of Geosciences, Beijing, China, in 2017. He is currently working as a Postdoctoral in a joint program of the Department of Computer Science and Technology, Tsinghua University, and Samsung Research China–Beijing (SRCB). His research interests include natural language processing and dialogue system, especially the tasks related to natural language understanding and natural language generation. He is the recipient of the Wuwenjun AI Award in 2019.



Guanyi Chen received the M.S. degree in artificial intelligence from the University of Edinburgh, Edinburgh, U.K., in 2016. He is currently working toward the Ph.D degree with the Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands. His research interests mainly lie in natural language generation, especially the computational modeling of human language production.



Minlie Huang (Member, IEEE) received the Ph.D. degree from Tsinghua University, Beijing, China, in 2006. He is currently an Associate Professor with the Department of Computer Science and Technology, Tsinghua University. He has authored or coauthored more than 60 papers in premier conferences and journals (ACL, EMNLP, AAAI, IJCAI, WWW, SIGIR, etc.) His research interests include natural language processing, particularly in dialog systems, reading comprehension, and sentiment analysis. His work on emotional chatting machines was reported by MIT Technology Review, the Guardian, Nvidia, and many other mass media. He serves as standing reviewer for ACL, Area Chairs for ACL 2020/2016, EMNLP 2019/2014/2011, and Senior PC members for AAAI 2017–2020 and IJCAI 2017–2020, and reviewers for TASLP, TKDE, TOIS, TPAMI, etc. He is a nominee of ACL 2019 best demo papers, the recipient of IJCAI 2018 distinguished paper award, CCL 2018 best demo award, NLPCC 2015 best paper award, Hanvon Youngth Innovation Award in 2018, and Wuwenjun AI Award in 2019. He was supported by a NSFC key project, several NSFC regular projects, and many IT companies.